



Cisco Nexus 6000 Series NX-OS FabricPath Configuration Guide, Release 7.x

January 29, 2014

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Cisco Nexus 6000 Series NX-OS FabricPath Configuration Guide
© 2014 Cisco Systems, Inc. All rights reserved.




| | |
|---|-----------|
| Preface | 15 |
| Overview | 11 |
| Information About FabricPath | 11 |
| Information About Conversational MAC Address Learning | 12 |
| Layer 3 Routing Considerations for FabricPath | 12 |
| Hot Standby Router Protocol Support | 13 |
| No Support for the bind-vrf Command | 13 |
| Peer Gateway Not Recommended on Aggregation Layer | 13 |
| Prerequisites for FabricPath | 13 |
| Guidelines and Limitations for FabricPath | 14 |
| Licensing Requirements for FabricPath | 14 |
| Configuring FabricPath Switching | 15 |
| Information About FabricPath Switching | 15 |
| FabricPath Encapsulation | 16 |
| Default IS-IS Behavior with FabricPath | 17 |
| Conversational MAC Address Learning | 18 |
| Switching Using FabricPath | 19 |
| Conflict Resolution and Optional FabricPath Tunings | 110 |
| Configuring FabricPath Switching | 110 |
| Installing the FabricPath Feature Set | 111 |
| Uninstalling the FabricPath Feature Set | 112 |
| Enabling the FabricPath Feature Set | 113 |
| Disabling the FabricPath Feature Set | 114 |
| Creating a FabricPath Topology | 115 |
| Changing a FabricPath Switch ID (Optional) | 116 |
| Configuring the FabricPath Timers (Optional) | 117 |
| Disabling FabricPath Graceful Merges (Optional) | 118 |
| Forcing the Links to Come Up (Optional) | 119 |
| Verifying the FabricPath Switching Configuration | 120 |
| Monitoring FabricPath Switching Statistics | 120 |
| Configuration Example for FabricPath Switching | 121 |

| | |
|--|------------|
| Default Settings for FabricPath Switching | 121 |
| Feature History for Configuring FabricPath Switching | 122 |
| Configuring FabricPath Interfaces | 123 |
| Information About FabricPath Interfaces | 123 |
| FabricPath Interfaces | 123 |
| STP and the FabricPath Network | 124 |
| vPC+ | 125 |
| Anycast HSRP | 127 |
| Guidelines and Limitations for FabricPath Interfaces | 127 |
| Configuring FabricPath Interfaces | 128 |
| Configuring FabricPath Interfaces | 128 |
| Configuring the STP Priority with Rapid PVST+ | 130 |
| Configuring the STP Priority with MST | 131 |
| Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch (Optional) | 133 |
| Configuring a vPC+ Switch ID | 134 |
| Configuring an Anycast HSRP Bundle | 135 |
| Verifying the FabricPath Interface Configuration | 137 |
| Monitoring FabricPath Interface Statistics | 138 |
| Configuration Example for FabricPath Interface | 138 |
| Feature History for Configuring FabricPath Interfaces | 140 |
| Configuring FabricPath Forwarding | 141 |
| Information About FabricPath Forwarding | 141 |
| FabricPath Forwarding Overview | 141 |
| FabricPath VLANs | 142 |
| Forwarding Known Unicast Packets Using ECMP | 143 |
| Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets | 144 |
| Configuring FabricPath Forwarding | 145 |
| Setting the VLAN Mode to FabricPath or CE | 145 |
| Mapping FabricPath VLANs to a Topology | 146 |
| Configuring FabricPath Unicast Load Balancing (Optional) | 148 |
| Verifying the FabricPath Configuration | 149 |
| Configuration Example for FabricPath Forwarding | 150 |
| Default Settings for Fabric Path Forwarding | 151 |
| Feature History for Configuring FabricPath | 151 |

153

| | |
|--|------------|
| Configuring Advanced FabricPath Features | 153 |
| Information About Advanced FabricPath Layer 2 IS-IS Configurations | 153 |
| Overload Bit | 154 |
| Setting Advanced FabricPath Layer 2 IS-IS Parameters | 154 |
| Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally | 154 |
| Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface | 159 |
| Clearing Advanced FabricPath Layer 2 IS-IS Counters | 163 |
| Multi-Destination Tree | 165 |
| Information About MDT | 165 |
| Guidelines and Limitations for MDT | 165 |
| Configuring MDT for Default Topology | 166 |
| Configuring Per-Topology MDT | 166 |
| Verifying the FabricPath Advanced Configurations | 167 |
| Feature History for Configuring FabricPath Advanced Features | 168 |
| Configuring FabricPath in Transit Mode | 11 |
| Information about FabricPath in Transit Mode | 11 |
| Guidelines and Limitations | 11 |
| Configuring FabricPath in Transit Mode | 11 |
| Example: Configuring FabricPath in Transit Mode | 12 |
| Feature History for Configuring FabricPath in Transit Mode | 13 |
| Configuration Limits for Cisco NX-OS FabricPath | 11 |



Preface

This preface contains the following sections:

- [Audience, page 5](#)
- [Organization, page 5](#)
- [Document Conventions, page 6](#)
- [Related Documentation, page 7](#)
- [Documentation Feedback, page 7](#)
- [Obtaining Documentation and Submitting a Service Request, page 7](#)

Audience

This publication is for experienced users who configure and maintain Cisco NX-OS switches.

Organization

This reference is organized as follows:

| Chapter and Title | Description |
|--|---|
| Chapter 1, “Overview” | Describes an overview of the FabricPath feature. |
| Chapter 2, “Configuring FabricPath Switching” | Describes FabricPath encapsulation, Layer 2 IS-IS, and MAC address conversational learning. Also describes how to change timers for FabricPath and resolve switch ID conflicts. |
| Chapter 3, “Configuring FabricPath Interfaces” | Describes FabricPath interfaces and VLANs as well as how the feature interacts with Classical Ethernet networks and Spanning Tree Protocol. |
| Chapter 4, “Configuring FabricPath Forwarding” | Describes how the FabricPath network uses multiple paths to forward unicast, unknown unicast, broadcast, and multicast frames. |

| Chapter and Title | Description |
|--|--|
| Chapter 5, “Configuring Advanced FabricPath Features” | Describes how to configure more advanced parameters for Layer 2 IS-IS. |
| Chapter A, “Configuration Limits for Cisco NX-OS FabricPath” | Describes the configuration limits for FabricPath. |

Document Conventions

Command descriptions use these conventions:

| Convention | Description |
|---------------|---|
| boldface font | Commands and keywords are in boldface. |
| italic font | Arguments for which you supply values are in italics. |
| { } | Elements in braces are required choices. |
| [] | Elements in square brackets are optional. |
| [x y z] | Optional alternative keywords are grouped in brackets and separated by vertical bars. |
| string | A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks. |

Screen examples use these conventions:

| | |
|-----------------------------|---|
| screen font | Terminal sessions and information that the switch displays are in screen font. |
| boldface screen font | Information you must enter is in boldface screen font. |
| <i>italic screen font</i> | Arguments for which you supply values are in italic screen font. |
| < > | Nonprinting characters, such as passwords, are in angle brackets. |
| [] | Default responses to system prompts are in square brackets. |
| !, # | An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line. |

This document uses the following conventions:



Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Tip

Means *the following information will help you solve a problem*.

Related Documentation

Documentation for the Cisco Nexus 6000 Series Switch is available at the following URL:

http://www.cisco.com/en/US/products/ps12806/tsd_products_support_series_home.html

The documentation set is divided into the following categories:

Release Notes

The release notes are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/prod_release_notes_list.html

Installation and Upgrade Guides

The installation and upgrade guides are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/prod_installation_guides_list.html

Command References

The command references are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/prod_command_reference_list.html

Technical References

The technical references are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/prod_technical_reference_list.html

Configuration Guides

The configuration guides are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/products_installation_and_configuration_guides_list.html

Error and System Messages

The system message reference guide is available at the following URL:

http://www.cisco.com/en/US/products/ps12806/products_system_message_guides_list.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus6k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.



Overview

This chapter contains the following sections:

- [Information About FabricPath, page 1-1](#)
- [Information About Conversational MAC Address Learning, page 1-2](#)
- [Layer 3 Routing Considerations for FabricPath, page 1-2](#)
- [Prerequisites for FabricPath, page 1-3](#)
- [Guidelines and Limitations for FabricPath, page 1-4](#)
- [Licensing Requirements for FabricPath, page 1-4](#)

Information About FabricPath

The FabricPath feature provides the following:

- Allows Layer 2 multipathing in the FabricPath network.
- Provides built-in loop prevention and mitigation with no need to use the Spanning Tree Protocol (STP).
- Provides a single control plane for unicast, unknown unicast, broadcast, and multicast traffic.
- Enhances mobility and virtualization in the FabricPath network.

The software randomly assigns a unique switch ID to each switch that is enabled with FabricPath.

When a frame enters the FabricPath network from a Classical Ethernet (CE) network, the ingress interfaces encapsulate the frame with a FabricPath header. The software builds paths, called trees, through the FabricPath network and assigns a forwarding tag (FTag) by flow to all the traffic in the FabricPath network. When the frame leaves the FabricPath network to go to a CE network, the egressing interface decapsulates the frame and leaves the regular CE header.

The FabricPath network uses the Layer 2 Intermediate System-to-Intermediate System (IS-IS) protocol to forward traffic in the network using the FabricPath headers. Layer 2 IS-IS is different from Layer 3 IS-IS; the two protocols work independently. Layer 2 IS-IS requires no configuration and becomes operational when you enable FabricPath on the switch. The frames carry the same FTag that is assigned at ingress throughout the FabricPath network, and Layer 2 IS-IS allows all switches to have the same view of all the trees built by the software.

For multi-direction traffic, such as unknown unicast, multicast, and broadcast, the tree path is used only to avoid packet looping. FTag ensures that all switches forward packets on the same tree across the network, although FabricPath supports multiple trees. Known unicast packets are forwarded throughout the network using equal cost multipath (ECMP) and FTag trees are not used for these packets. Using ECMP and the trees, the software automatically load balances traffic throughout the FabricPath network.

**Note**

Precision Time Protocol (PTP) over FabricPath is not supported.

Information About Conversational MAC Address Learning

You can use conversational MAC address learning or traditional MAC address learning.

To use conversational MAC address learning, you must do the following:

- Enable FabricPath.
- Ensure VLANs do not have switch virtual interface (SVI) enabled.

Conversational MAC address learning means that each switch learns only those MAC addresses for interested hosts, rather than all MAC addresses in the domain. Each switch learns only those MAC addresses that are actively speaking with it. In this way, conversational MAC learning consists of a three-way handshake.

This selective learning, or conversational MAC address learning, allows you to scale the network beyond the limits of individual switch MAC address tables.

All FabricPath VLANs use conversational MAC address learning.

**Note**

CE VLANs support only traditional MAC address learning, where each switch learns the MAC addresses of all hosts in the network.

Layer 3 Routing Considerations for FabricPath

This section provides information you need to consider when you implement routing technologies for Layer 3 routing that take advantage of FabricPath to improve both unicast and multicast routing in your data center.

This section includes the following topics:

- [Hot Standby Router Protocol Support, page 1-3](#)
- [No Support for the bind-vrf Command, page 1-3](#)

In addition to the above, FabricPath includes support for the **delay restore interface-vlan** command and the **ip arp synchronize** command. For more information about these commands, see the command reference for your platform.

Hot Standby Router Protocol Support

You can configure Hot Standby Router Protocol (HSRP) in Active/Active mode between pairs of switches in the vPC+ domains and access layer, and between pairs of switches in the aggregation layer.

If you configure a peer link between a pair of switches in the aggregation layer, HSRP runs in Active/Active mode for those switches. If you do not configure a peer link between the switches in the aggregation layer, HSRP runs in Active/Standby mode for those switches.

**Note**

If a data center outage occurs and you enable HSRP before the vPC+ successfully comes up, traffic loss can occur. You must enable an HSRP delay to give the vPC time to stabilize. If you enable both an HSRP delay and a preemption delay, the Cisco Nexus 6000 Series switches allows Layer 2 switching only after both timers expire.

The delay option is available only with HSRP. If you use any other FHRP, traffic loss is still possible.

No Support for the bind-vrf Command

One consequence of the single-DR implementation for FabricPath is that you do not need to bind a virtual routing and forwarding (VRF) instance to a vPC. As a result, the **vpc bind-vrf** command is not supported for FabricPath.

If you want to configure an existing vPC domain for FabricPath, you must first use the **no vpc bind-vrf** command to remove the static binding between the vPC and the VRF.

Peer Gateway Not Recommended on Aggregation Layer

We do not recommend that you configure a peer gateway on the aggregation layer, as Fabric Path does not require one.

The **peer-gateway** command is supported by vPC and vPC+ at the access. If traffic from a vPC for peer-1's MAC address is sent to peer-2, the peer gateway configuration ensures that peer-2 can route the packet on peer-1's behalf. However, a FabricPath topology does not have a port channel that can cause traffic for one peer switch to be sent to the other peer switch. Therefore, FabricPath does not require peer gateway configuration, even though there is a peer link provisioned for HSRP Active-Active support.

Prerequisites for FabricPath

FabricPath has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must be logged onto the switch.
- You must ensure an Enhanced Layer 2 license is installed on the switch.
- If you want to configure an existing vPC domain for FabricPath, you must first use the **no vpc bind-vrf** command to remove the static binding between the vPC and the VRF.

Guidelines and Limitations for FabricPath

FabricPath has the following configuration guidelines and limitations:

- FabricPath switches and interfaces carry only FabricPath-encapsulated traffic.
- You must install and enable FabricPath on each switch before you can view or access the commands. See the “[Configuring FabricPath Switching](#)” section on page 2-10 for more information.
- STP does not run inside a FabricPath network.
- The following guidelines apply to private VLAN configurations when you are running FabricPath:
 - All VLANs in a private VLAN must be in the same VLAN mode: either CE or FabricPath. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the private VLAN. The software remembers the configurations, and if you change the VLAN mode later, that VLAN becomes active in the specified private VLAN.
 - FabricPath ports cannot be put into a private VLAN.
- FabricPath does not support hierarchical static MAC addresses. That is, you cannot configure static FabricPath ODAs or OSAs; you can only configure CE static MAC addresses.
- FabricPath does not support VTP when in the same VDC. You must disable VTP when the FabricPath feature set is enabled on the VDC.

**Note**

For information about FabricPath and DHCP snooping, including guidelines and limitations, see the *Cisco Nexus 6000 Series NX-OS Security Configuration Guide, Release 6.0*.

Licensing Requirements for FabricPath

FabricPath requires an Enhanced Layer 2 license. You must install this license on every switch in a FabricPath network.

In addition, if you want to take advantage of the Layer 3 routing interactions with FabricPath, you must install a LAN Base Services license on all switches with a Layer 3 card.



Configuring FabricPath Switching

This chapter contains the following sections:

- [Information About FabricPath Switching, page 2-5](#)
- [Configuring FabricPath Switching, page 2-10](#)
- [Verifying the FabricPath Switching Configuration, page 2-20](#)
- [Monitoring FabricPath Switching Statistics, page 2-20](#)
- [Configuration Example for FabricPath Switching, page 2-21](#)
- [Default Settings for FabricPath Switching, page 2-21](#)
- [Feature History for Configuring FabricPath Switching, page 2-22](#)



Note

For information about prerequisites, guidelines and limitations, and licensing requirements for FabricPath, see [Chapter 1, “Overview.”](#)

Information About FabricPath Switching

FabricPath switching allows multipath networking at the Layer 2 level. The FabricPath network still delivers packets on a best-effort basis (which is similar to the Classical Ethernet [CE] network), but the FabricPath network can use multiple paths for Layer 2 traffic. In a FabricPath network, you do not need to run the Spanning Tree Protocol (STP). Instead, you can use FabricPath across data centers, some of which have only Layer 2 connectivity, with no need for Layer 3 connectivity and IP configurations.

The FabricPath encapsulation facilitates MAC mobility and server virtualization, which means that you can physically move the Layer 2 node but retain the same MAC address and VLAN association for the virtual machine. FabricPath also allows LAN extensions across data centers at Layer 2, which is useful in disaster recovery operations, as well as clustering applications such as databases. Finally, FabricPath is very useful in high-performance, low-latency computing.

With FabricPath, you use the Layer 2 intermediate System-to-Intermediate System (IS-IS) protocol for a single control plane that functions for unicast, broadcast, and multicast packets. There is no need to run the Spanning Tree Protocol (STP); it is a purely Layer 2 domain. This FabricPath Layer 2 IS-IS is a separate process from Layer 3 IS-IS.

Cisco supports the conversation-based MAC learning schema. Conversational learning can only be applied to FabricPath VLANs. Using FabricPath and conversational MAC address learning, the switch has to learn far fewer MAC addresses, which results in smaller, more manageable MAC tables.

This section includes the following topics:

- [FabricPath Encapsulation, page 2-6](#)
- [Default IS-IS Behavior with FabricPath, page 2-7](#)
- [Conversational MAC Address Learning, page 2-8](#)
- [Switching Using FabricPath, page 2-9](#)
- [Conflict Resolution and Optional FabricPath Tunings, page 2-10](#)

FabricPath Encapsulation

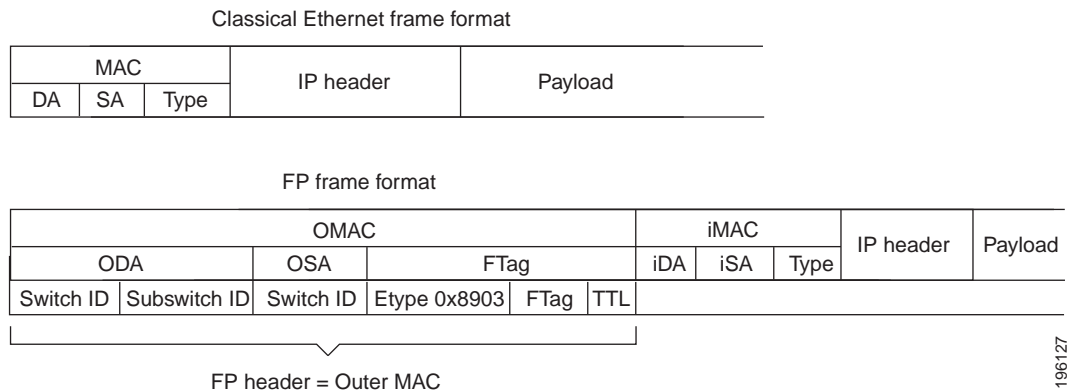
This section includes the following topics:

- [FabricPath Headers, page 2-6](#)
- [Forwarding Tags \(FTags\), page 2-7](#)

FabricPath Headers

When a frame enters the FabricPath network, the software encapsulates the Layer 2 frame with a new FabricPath header. The switch IDs that the software assigns to each FabricPath switch as it enters the FabricPath network is used as the outer MAC destination address (ODA) and outer MAC source address (OSA) in the FabricPath header. [Figure 2-1](#) shows the FabricPath header encapsulation of the classical Ethernet (CE) frame.

Figure 2-1 FabricPath Frame Encapsulation



The software applies the encapsulation on the ingress edge port of the FabricPath network and decapsulates the frame on the egress edge port of the FabricPath network; all the ports within the FabricPath network are FabricPath ports that use only the hierarchical MAC address (see [Chapter 3, “Configuring FabricPath Interfaces”](#), for more information on configuring FabricPath interfaces). This feature greatly reduces the size of the MAC tables in the core of the FabricPath network.

The software automatically assigns each switch in the FabricPath network with a unique switch ID. Optionally, you can configure the switch ID for the FabricPath switch.

The outer source address (OSA) is the FabricPath switch ID of the switch where the frame ingresses the FabricPath network, and the outer destination address (ODA) is the FabricPath switch ID of the switch where the frame egresses the FabricPath network. When the frame egresses the FabricPath network, the FabricPath switch strips the FabricPath header, and the original CE frame continues on the CE network.

The FabricPath network uses only the OSA and ODA, with the Layer 2 IS-IS protocol transmitting the topology information. Both the FabricPath ODA and OSA are in a standard MAC format (xxxx.xxxx.xxxx).

The FabricPath hierarchical MAC address carries the reserved EtherType 0x8903.

When the frame is originally encapsulated, the software sets the time to live (TTL) to 32; on each hop through the FabricPath network, each switch decrements the TTL by 1. If the TTL reaches 0, that frame is discarded. This feature prevents the continuation of any loops that might form in the network.

Forwarding Tags (FTags)

The Forwarding Tag (FTag) in the FabricPath header specifies which path the packet traverses throughout the FabricPath network. The switch uses the FTag-specified paths for multidestination packets that enter the FabricPath network. The FTag, which is a fixed route that the switch learns from the topology, is a 10-bit field with the values from 1 to 1024 (see [Chapter 4, “Configuring FabricPath Forwarding,”](#) for more information).

The software assigns the FTag to the FabricPath header on the edge port as the frame ingresses the FabricPath network. All subsequent FabricPath switches in that FabricPath network recognize the FTag. Each FTag is unique within the FabricPath topology.

Default IS-IS Behavior with FabricPath

The interfaces in a FabricPath network run only the FabricPath Layer 2 IS-IS protocol; you do not need to run STP in the FabricPath network because FabricPath Layer 2 IS-IS discovers topology information dynamically.



Note

Cisco Nexus 6000 Series switches supports two Fabricpath topologies, the default or base topology (topology 0), and another optional topology (for example, topology 1).

FabricPath Layer 2 IS-IS is a dynamic link-state routing protocol that detects changes in the network topology and calculates loop-free paths to other nodes in the network. Each FabricPath switch maintains a link-state database (LSDB) that describes the state of the network; each switch updates the status of the links that are adjacent to the switch. The FabricPath switch sends advertisements and updates to the LSDB through all the existing adjacencies. FabricPath Layer 2 IS-IS protocol packets do not conflict with standard Layer 2 IS-IS packets because the FabricPath packets go to a different Layer 2 destination MAC address than that used by standard IS-IS for IPv4/IPv6 address families.

The software sends hello packets on the FabricPath core ports to form adjacencies. After the software forms IS-IS adjacencies, the FabricPath unicast traffic uses the equal cost multipath (ECMP) feature of Layer 2 IS-IS to forward traffic. ECMP provides up to 16 paths for unicast traffic.

Within the FabricPath network, you use a single control plane protocol, Layer 2 IS-IS, for all unicast, multicast, and broadcast traffic. To use the basic FabricPath functionality, you do not need to configure Layer 2 IS-IS. The control plane Layer 2 IS-IS comes up and runs automatically when you enable FabricPath on the switch.

The loop-free Layer 2 IS-IS protocol builds two trees for the topology. Multicast and unknown unicast traffic is load-balanced between the two trees (see [Chapter 4, “Configuring FabricPath Forwarding,”](#) for more information about trees and topology).

FabricPath Layer 2 IS-IS is based on the standard IS-IS protocol with the following extensions for the FabricPath environment:

- FabricPath has a single IS-IS area with no hierarchical Layer 1/Layer 2 routing as prescribed within the IS-IS standard. All switches within the FabricPath network are in a single Layer 1 area.
- The switch uses a MAC address that is different from the MAC address used for Layer 3 IS-IS instances.
- The software adds a new sub-TLV that carries switch ID information, which is not in standard IS-IS. This feature allows Layer 2 information to be exchanged through the existing IS-IS protocol implementation.
- Each switch computes its shortest path to every other switch in the network using the shortest-path first (SPF) algorithm. This path is used for forwarding unicast FabricPath frames. FabricPath Layer 2 IS-IS uses the standard IS-IS functionality to populate up to 16 routes for a given destination switch. The software uses multiple equal-cost available parallel links that provide ECMP.
- FabricPath IS-IS introduces certain modifications to the standard IS-IS in order to support the construction of broadcast and multicast trees (identified by the FTags). Specifically, using FabricPath, the software constructs two loop-free trees for forwarding multideestination traffic.

Once the adjacency is established among the switches in the FabricPath network, the software sends update information to all neighbors.

By default, you can run Layer 2 IS-IS with FabricPath with the default configuration. However, you can fine-tune some of the Layer 2 IS-IS parameters (see [Chapter 5, “Configuring Advanced FabricPath Features,”](#) for information on configuring optional IS-IS parameters).

Additionally, FabricPath IS-IS helps to ensure that each switch ID in steady-state is unique within the FabricPath network. If FabricPath networks merge, switch IDs might collide. If the IDs are all dynamically assigned, FabricPath IS-IS ensures that this conflict is resolved without affecting any FabricPath traffic in either network.

Conversational MAC Address Learning

In traditional MAC address learning, each switch learns the MAC address of every other host in the network. When you configure a VLAN for conversational learning, the associated switches learn only those MAC addresses that are actively in conversation with the hosts connected to the edge interfaces of those switches. Not all switches have to learn all the MAC addresses, which greatly reduces the size of the MAC address tables.

You can optimize the MAC learning process. Conversational MAC learning is configured per VLAN. All FabricPath VLANs use conversational learning if a switch virtual interface (SVI) is not enabled on the switch. (See [Chapter 4, “Configuring FabricPath Forwarding,”](#) for more information on CE and FabricPath VLANs.)

Conversational MAC address learning results in MAC address tables that are much smaller for FabricPath.

The MAC address learning modes are the traditional learning and conversational learning. The following VLAN modes have the following MAC learning modes:


- FabricPath VLANs—Conversational MAC learning is enabled unless you enable an SVI.
- CE VLANs—Only traditional learning.

With conversational MAC learning, the switch learns the source MAC address of an ingressing frame only if destination MAC address is already learnt and present in the MAC address table. It requires three-way handshake for the MAC address to be learnt. Unknown MAC address are forwarded, or flooded throughout the network.

Switching Using FabricPath

The FabricPath hierarchical MAC address scheme and conversational learning result in much smaller, conversational learning MAC tables within the FabricPath network. Within the FabricPath network, the software uses Layer 2 IS-IS to transmit topology information. The switches on the edge of the network, which use conversational MAC address learning, do not have to learn all the MAC addresses in the network (see [Figure 2-2](#)).

Figure 2-2 FabricPath Ports Use Only the FabricPath Header to Switch Frames



MAC mobility is expedited using the FabricPath hierarchical MAC addresses. That is, when you want to move a host and keep its same MAC address and VLANs, only the switches at the edge of the FabricPath network track this change. Within the FabricPath network, the FabricPath switches update their tables with only the outer MAC addresses (ODA and OSA) that have changed from the FabricPath encapsulation.

The switch on the edge of the FabricPath network encapsulates the original frame inside the FabricPath header. Once the frame reaches the last, or directly connected, FabricPath switch, the egress interface strips the FabricPath header and forwards the frame as a normal CE frame.

The ports at the edge of a FabricPath network can use conversational learning to learn only those MAC addresses that the specified edge port is having a bidirectional conversation with. Every edge switch does not have to learn the MAC address of every other edge switch; they only need to learn the MAC addresses of the speakers.

As the frame traverses the FabricPath network, all the switches work only with the FabricPath header. So, the FabricPath switches work only with the ODAs and OSAs; they do not need to learn the MAC address for any of the CE hosts or other switches attached to the network. The hierarchical MAC addressing provided by the FabricPath headers results in much smaller MAC tables in the FabricPath network. The switches in the FabricPath network only need to know how to forward frames to another FabricPath switch so they can forward traffic without requiring large MAC address lookup tables in the core of the network.

The switches in the FabricPath network decrement the TTL in the FabricPath header by 1 at each hop. When the TTL reaches 0, the packet is dropped. This process prevents the continuation of any loops that might form in the network.

Conflict Resolution and Optional FabricPath Tunings

After you enable FabricPath in all switches, the software automatically assigns a random switch ID to each FabricPath switch. The switch ID is a 12-bit value that is dynamically assigned to every switch in the FabricPath network, with each switch being a unique value in that FabricPath network. Optionally, you can configure a specific switch ID. If any of the switch IDs in the FabricPath network are not unique, the software provides automatic conflict resolution.

The software chooses a random value for the switch ID and sets this value as tentative during a period when the software waits to hear if this value is already in use. If this value is being used by another switch in the network, the software begins a conflict resolution process. The switch with the lower switch ID keeps the specified value and the other switch gets a new value for its switch ID.

If a single switch joins an existing FabricPath network, that switch changes its switch ID value rather than any of the existing switches in the network changing values. If the specified value is not in use by another switch or after the conflict is resolved, the switch ID is marked as confirmed.

Graceful migration provides that there is no traffic disruption if a conflict arises in the resources, such as two switches that temporarily have the same switch ID.

The FabricPath interfaces come up, but they are not operational until the switch checks for FabricPath conflicts and resolves those conflicts.

The FabricPath resources timers have default values, but you can also change the timer values. You can tune the switch to wait longer or shorter periods to check the conflicts.

Some of the important processes of the FabricPath network are as follows:

- Achieves a conflict-free allocation of switch IDs and FTags
- Provides graceful resource migration during network merges or partition healing
- Supports static switch IDs
- Provides fast convergence during link bringup or network merges

FabricPath uses the Layer 2 IS-IS protocol to transport the database to all switches in the network. The information is distributed among the FabricPath network switches using an IS-IS TLV. Each switch sends its version of the database that contains information about all the switches.

The software allocates the FabricPath values, guarantees their uniqueness within the FabricPath network, and deletes the value from the database once that resource is no longer needed.

When you manually configure static switch IDs for the switch, the automatic conflict resolution process does not work and the network does not come up. You will see syslog messages about the conflict and must manually change one or more switch IDs of the switches in the network.

Configuring FabricPath Switching

After you enable FabricPath switching on each switch, the encapsulation, default IS-IS, and learning occur automatically.

**Note**

You must install the FabricPath feature set before you enable FabricPath on the switch. See the [“Enabling the FabricPath Feature Set” section on page 2-13](#) for complete information on installing and enabling the FabricPath feature set.

Instead of using the default values, you can optionally configure the following FabricPath values that the software uses for conflict resolution and other tunings:

- Switch ID for the switch that is used globally in the FabricPath network
- Timers
- Graceful merge of FabricPath networks (Enabled by default. You might experience traffic drops if the feature is disabled.)
- A one-time forcing of the links to come up

This section includes the following topics:

- [Installing the FabricPath Feature Set, page 2-11](#)
- [Uninstalling the FabricPath Feature Set, page 2-12](#)
- [Enabling the FabricPath Feature Set, page 2-13](#)
- [Disabling the FabricPath Feature Set, page 2-14](#)
- [Creating a FabricPath Topology, page 2-15](#)
- [Configuring the FabricPath Timers \(Optional\), page 2-17](#)
- [Disabling FabricPath Graceful Merges \(Optional\), page 2-18](#)
- [Forcing the Links to Come Up \(Optional\), page 2-19](#)

Installing the FabricPath Feature Set

You must install the FabricPath feature set on all switches on which you want to configure FabricPath.

BEFORE YOU BEGIN

Ensure that you have installed an Enhanced Layer 2 license on all switches on which you want to install FabricPath.

SUMMARY STEPS

1. **configure terminal**
2. **install feature-set fabricpath**
3. **exit**
4. (Optional) **show feature-set**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>install feature-set fabricpath</code> Example: switch(config)# <code>install feature-set fabricpath</code> | Installs the FabricPath feature set on the switch. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch# | Exits configuration mode. |
| Step 4 | <code>show feature-set</code> Example: switch# <code>show feature-set</code> | (Optional) Displays which feature sets are enabled on the switch. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to install the FabricPath feature set on a switch:

```
switch# configure terminal
switch(config)# install feature-set fabricpath
switch(config)#
```

Uninstalling the FabricPath Feature Set

BEFORE YOU BEGIN

Ensure that you have disabled the FabricPath featureset on all switches from which you want to uninstall FabricPath. See the [“Disabling the FabricPath Feature Set”](#) section on page 2-14 for more information.

SUMMARY STEPS

1. **configure terminal**
2. **no install feature-set fabricpath**
3. **exit**
4. (Optional) **show feature-set**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>no install feature-set fabricpath</code> Example: switch(config)# <code>no feature-set fabricpath</code> | Uninstalls the FabricPath feature set from the switch. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch# | Exits configuration mode. |
| Step 4 | <code>show feature-set</code> Example: switch# <code>show feature-set</code> | (Optional) Displays which feature sets are enabled on the switch. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to uninstall the FabricPath feature set from a switch:

```
switch# configure terminal
switch(config)# no install feature-set fabricpath
switch(config)#
```

Enabling the FabricPath Feature Set

You must enable the FabricPath feature set on all switches on which you want to run FabricPath. You cannot access the commands that you use to configure the feature until the feature set is installed and enabled.

BEFORE YOU BEGIN

Ensure that you have installed the FabricPath feature on the switch.

SUMMARY STEPS

1. `configure terminal`
2. `feature-set fabricpath`
3. `exit`
4. (Optional) `show feature-set`
5. (Optional) `copy running-config startup-config`

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>feature-set fabricpath</code> Example: switch(config)# <code>feature-set fabricpath</code> | Enables the FabricPath feature set on the switch. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch# | Exits configuration mode. |
| Step 4 | <code>show feature-set</code> Example: switch# <code>show feature-set</code> | (Optional) Displays which feature sets are enabled on the switch. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to enable the FabricPath feature set on a switch:

```
switch# configure terminal
switch(config)# feature-set fabricpath
switch(config)#
```

Disabling the FabricPath Feature Set

When you disable the FabricPath feature set, the switch clears all FabricPath configurations from the switch and you can no longer see any of the CLI commands that you need to configure FabricPath.



Note

If your FabricPath configuration is large (multiple megabytes in size), disabling the FabricPath feature set may take some time to complete.

SUMMARY STEPS

1. `configure terminal`
2. `no feature-set fabricpath`
3. `exit`
4. (Optional) `show feature-set`
5. (Optional) `copy running-config startup-config`

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>no feature-set fabricpath</code> Example: switch(config)# <code>no feature-set fabricpath</code> | Disables the FabricPath feature in the default VDC. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch# | Exits configuration mode. |
| Step 4 | <code>show feature-set</code> Example: switch# <code>show feature-set</code> | (Optional) Displays which feature sets are enabled on the switch. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to disable the FabricPath feature:

```
switch# configure terminal
switch(config)# no feature-set fabricpath
switch(config)#
```

Creating a FabricPath Topology

You must install the FabricPath feature set on all switches on which you want to configure FabricPath.

SUMMARY STEPS

1. `configure terminal`
2. `fabricpath topology topology-number`
3. (Optional) `copy running-config startup-config`

DETAILED STEPS

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>fabricpath topology topology-number</code> Example: switch(config)# <code>fabricpath topology 1</code> | Specifies the FabricPath topology number. Note Cisco Nexus Series only supports two topologies; the default or base topology (topology 0) and another optional topology (for example, topology 1). |
| Step 3 | <code>show fabricpath topology</code> Example: switch# <code>show fabricpath</code> | (Optional) Displays information about the FabricPath IS-IS topology. |
| Step 4 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure the FabricPath topology:

```
switch# configure terminal
switch(config)# fabricpath topology 1
switch(config)#
```

Changing a FabricPath Switch ID (Optional)

By default, FabricPath assigns a unique switch ID to a switch after you enable FabricPath on the switch. However, you can manually change the switch ID if you prefer.



Note You do not lose any traffic during switch ID changes.

SUMMARY STEPS

1. `configure terminal`
2. `fabricpath switch-id value`
3. `exit`
4. (Optional) `show fabricpath switch-id`
5. (Optional) `copy running-config startup-config`

DETAILED STEPS

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>fabricpath switch-id value</code> Example: switch(config)# <code>fabricpath switch-id 25</code> switch(config)# | Specifies the switch ID. The range is from 1 to 4094. There is no default value. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch)# | Exits the configuration mode. |
| Step 4 | <code>show fabricpath switch-id</code> Example: switch# <code>show fabricpath switch-id</code> | (Optional) Displays information on the switch IDs. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to manually change the FabricPath switch ID on a switch to 25:

```
switch# configure terminal
switch(config)# fabricpath switch-id 25
switch(config)#
```

Configuring the FabricPath Timers (Optional)

You can configure the following FabricPath timers:

- `allocate-delay`—Configures the delay for a new switch ID to be propagated throughout the network before that value becomes available and permanent.
- `linkup-delay`—Configures the link bringup delay to detect conflicts in the switch ID. If the software does find a conflict, the software takes some time to resolve the conflict and bring FabricPath to an operational state.
- `transition-delay`—Configures the delay for propagating a transitioned value in the network; during this period, all old and new switch ID values exist in the network. This situation occurs only while the link comes up and the software checks to see if the network has two identical switch IDs.

SUMMARY STEPS

1. `configure terminal`
2. `fabricpath timers {allocate-delay seconds | linkup-delay seconds | transition-delay seconds}`
3. `exit`
4. (Optional) `show fabricpath timers`

5. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|--|--|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters global configuration mode. |
| Step 2 | fabricpath timers {allocate-delay seconds / linkup-delay seconds / transition-delay seconds} Example: switch(config)# fabricpath allocate-delay 600 switch(config)# | Specifies the FabricPath timer values. The range is from 1 to 1200 seconds for each of these timers. The default values are as follows: <ul style="list-style-type: none"> • allocate-delay—10 seconds • linkup-delay—10 seconds • transition-delay—10 seconds |
| Step 3 | exit Example: switch(config)# exit switch)# | Exits the configuration mode. |
| Step 4 | show fabricpath timers Example: switch# show fabricpath timers | (Optional) Displays information on the FabricPath timers. |
| Step 5 | copy running-config startup-config Example: switch# copy running-config startup-config | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure the allocation-delay for a FabricPath timer to 600 seconds:

```
switch# configure terminal
switch(config)# fabricpath timer allocate-delay 600
switch(config)#
```

Disabling FabricPath Graceful Merges (Optional)

By default, graceful merges are enabled. If desired, you can disable this aspect of the FabricPath feature. However, if you disable graceful merges, you must do so on all FabricPath switches.



Note

You may experience traffic drops if you disable this feature.

SUMMARY STEPS

1. **configure terminal**
2. **[no] fabricpath graceful-merge disable**
3. **exit**
4. (Optional) **show running-config**

5. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>[no] fabricpath graceful-merge disable</code> Example: switch(config)# <code>fabricpath graceful-merge disable</code> switch(config)# | Disables graceful merge of the FabricPath feature. To reenale this feature, enter the no form of the command. |
| Step 3 | <code>exit</code> Example: switch(config)# <code>exit</code> switch)# | Exits configuration mode. |
| Step 4 | <code>show running-config</code> Example: switch# <code>show running-config</code> | (Optional) Displays information on the configuration running on the switch. |
| Step 5 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to disable graceful merges on a FabricPath switch:

```
switch# configure terminal
switch(config)# fabricpath graceful-merge disable
switch(config)#
```

Forcing the Links to Come Up (Optional)



Note

We do NOT recommend that you use the **fabricpath force link-bringup** command to force the network links to connect.

As a one-time event, you can force the FabricPath network links to connect if they are not coming up because of switch ID conflicts or other problems in the network.

You must repeat this configuration on every switch in the FabricPath network.

SUMMARY STEPS

Step 1 **fabricpath force link-bringup**

DETAILED STEPS

| | Command | Purpose |
|--------|--|---|
| Step 1 | <pre>fabricpath force link-bringup</pre> <p>Example:</p> <pre>switch# fabricpath force link-bringup switch#</pre> | <p>Forces the FabricPath network links to come up as a one-time event.</p> <p>Note This command is not saved when you enter the copy running-config startup-config command.</p> |

This example shows how to force the FabricPath network links to come up as a one-time event:

```
switch# fabricpath force link-bringup
switch#
```

Verifying the FabricPath Switching Configuration

To display FabricPath switching information, perform one of the following tasks:

| Command | Purpose |
|--|---|
| show feature-set | Displays whether FabricPath is enabled or not. |
| show fabricpath conflict {all [<i>detail</i>] link [<i>detail</i>] switch-id [<i>detail</i>] transitions [<i>detail</i>]} | Displays information on conflicts in the FabricPath network. |
| show fabricpath switch-id [<i>local</i>] | Displays information on the FabricPath network by switch ID. |
| show fabricpath system-id { <i>mac-addr</i> } | Displays information on the FabricPath network by system ID. |
| show fabricpath timers | Displays settings for the allocate-delay, linkup-delay, and transition-delay timers for the FabricPath network. |

See [Chapter 5, “Configuring Advanced FabricPath Features,”](#) for more commands that display the FabricPath switching functionality.

Monitoring FabricPath Switching Statistics

Use the following commands to display or clear the FabricPath switching statistics:

- **load- interval** {interval *seconds* {1 | 2 | 3}}
- **show interface counters** [*module module*]
- **show interface counters detailed** [*all*]
- **show interface counters errors** [*module module*]
- **clear counters** [*interface*]

Configuration Example for FabricPath Switching

Step 1 Install FabricPath on all the switches.

```
switch# configure terminal
switch(config)# install feature-set fabricpath
switch(config)#
```

Step 2 Enable FabricPath on all the switches.

```
switch# configure terminal
switch(config)# feature-set fabricpath
switch(config)#
```

Step 3 (Optional) Manually configure a switch ID for the FabricPath switch.

```
switch# configure terminal
switch(config)# fabricpath switch-id 25
switch(config)#
```

Step 4 Save the configuration.

```
switch(config)# copy running-config startup-config
switch(config)#
```

Default Settings for FabricPath Switching

Table 2-1 lists the default settings for FabricPath switching parameters.

Table 2-1 Default FabricPath Parameters

| Parameters | Default |
|---------------------------|--|
| FabricPath | Disabled |
| MAC address learning mode | <ul style="list-style-type: none"> FabricPath VLANs—Only conversational learning CE VLANs—Traditional (nonconversational) learning |
| allocate-delay timer | 10 seconds |
| linkup-delay timer | 10 seconds |
| transition-delay timer | 10 seconds |
| graceful merge | Enabled |

Feature History for Configuring FabricPath Switching

Table 2-2 lists the release history for these features.

Table 2-2 Feature History for FabricPath

| Feature Name | Releases | Feature Information |
|----------------------|-------------|------------------------------|
| FabricPath Switching | 6.0(2)N1(1) | This feature was introduced. |



Configuring FabricPath Interfaces

This chapter contains the following sections:

- [Information About FabricPath Interfaces, page 3-23](#)
- [Configuring FabricPath Interfaces, page 3-28](#)
- [Verifying the FabricPath Interface Configuration, page 3-37](#)
- [Monitoring FabricPath Interface Statistics, page 3-38](#)
- [Configuration Example for FabricPath Interface, page 3-38](#)
- [Feature History for Configuring FabricPath Interfaces, page 3-40](#)



Note

For information about prerequisites, guidelines and limitations, and licensing requirements for FabricPath, see [Chapter 1, “Overview.”](#)

Information About FabricPath Interfaces

This section includes the following sections:

- [FabricPath Interfaces, page 3-23](#)
- [STP and the FabricPath Network, page 3-24](#)
- [vPC+, page 3-25](#)
- [Anycast HSRP, page 3-27](#)

FabricPath Interfaces

After you enable FabricPath on the switches that you are using, you can configure an Ethernet interface or a port-channel interface as a FabricPath interface. If one member of the port channel is in FabricPath mode, all the other members will be in FabricPath mode. After you configure the interface as a FabricPath interface, it automatically becomes a trunk port, capable of carrying traffic for multiple VLANs. You can also configure all the ports on the switch as FabricPath interfaces simultaneously.

The following interface modes carry traffic for the following types of VLANs:

- Interfaces on the switch that are configured as FabricPath interfaces can carry traffic only for FabricPath VLANs.

- Interfaces on the switch that are not configured as FabricPath interfaces carry traffic for the following:
 - FabricPath VLANs
 - Classical Ethernet (CE) VLANs

See [Chapter 4, “Configuring FabricPath Forwarding,”](#) for information about FabricPath and CE VLANs.

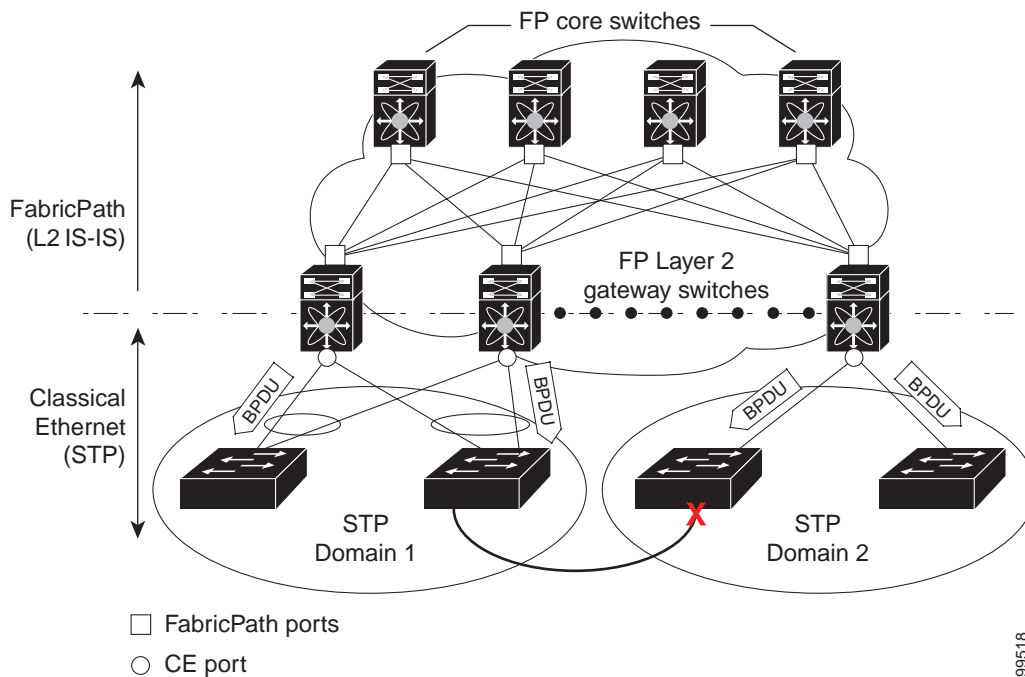
The FabricPath interfaces connect only to other FabricPath interfaces within the FabricPath network. These FabricPath ports operate on the information in the FabricPath headers and Layer 2 Intermediate System-to-Intermediate System (IS-IS) only, and they do not run STP. These ports are aware only of FabricPath VLANs; they are unaware of any CE VLANs. By default, all VLANs are allowed on a trunk port, so the FabricPath interface carries traffic for all FabricPath VLANs.

STP and the FabricPath Network

The Layer 2 gateway switches, which are on the edge between the CE and the FabricPath network, must be configured as the root for all Spanning Tree Protocol (STP) domains that are connected to a FabricPath network.

The STP domains do not cross into the FabricPath network (see [Figure 3-1](#)).

Figure 3-1 STP Boundary Termination at FabricPath Network Border



You must configure the FabricPath Layer 2 gateway switch to have the lowest STP priority of all the switches in the STP domain to which it is attached. You must also configure all the FabricPath Layer 2 gateway switches that are connected to one FabricPath network to have the *same* priority. The software assigns the bridge ID for the Layer 2 gateway switches from a pool of reserved MAC addresses.

To ensure a loop-free topology for the CE/FabricPath hybrid network, the FabricPath network automatically displays as a single bridge to all connected CE switches.

**Note**

You must set the STP priority on all FabricPath Layer 2 gateway switches to a value low enough to ensure that they become root for any attached STP domains.

Other than configuring the STP priority on the FabricPath Layer 2 gateway switches, you do not need to configure anything for the STP to work seamlessly with the FabricPath network. Only connected CE switches form a single STP domain. Those CE switches that are not interconnected form separate STP domains (see [Figure 3-1](#)).

All CE interfaces should be designated ports, which occurs automatically, or they are pruned from the active STP topology. If the software prunes any port, the software returns a syslog message. The software clears the port again only when that port is no longer receiving superior BPDUs.

The FabricPath Layer 2 gateway switch also propagates the Topology Change Notifications (TCNs) on all its CE interfaces.

The FabricPath Layer 2 gateway switches terminate STP. The set of FabricPath Layer 2 gateway switches that are connected by STP forms the STP domain. Because many FabricPath Layer 2 gateway switches might be attached to a single FabricPath network, there might also be many separate STP domains (see [Figure 3-1](#)). The switches in the separate STP domains need to know the TCN information only for the domain to which they belong. You can configure a unique STP domain ID for each separate STP domain that connects to the same FabricPath network. The Layer 2 IS-IS messages carry the TCNs across the FabricPath network. Only those FabricPath Layer 2 gateway switches in the same STP domain as the TCN message need to act and propagate the message to connected CE switches.

When a FabricPath Layer 2 gateway switch receives a TCN for the STP domain it is part of, it takes the following actions:

- Flushes all remote MAC addresses for that STP domain and the MAC addresses on the designated port.
- Propagates the TCN to the other switches in the specified STP domain.

The switches in the separate STP domains need to receive the TCN information and then flush all remote MAC addresses reachable by the STP domain that generated the TCN information.

vPC+

A virtual port channel+ (vPC+) domain allows a classical Ethernet (CE) vPC domain and a Cisco FabricPath cloud to interoperate. A vPC+ also provides a First Hop Routing Protocol (FHRP) active-active capability at the FabricPath to Layer 3 boundary.

**Note**

vPC+ is an extension to virtual port channels (vPCs) that run CE only.. You cannot configure a vPC+ domain and a vPC domain on the same Cisco Nexus device.

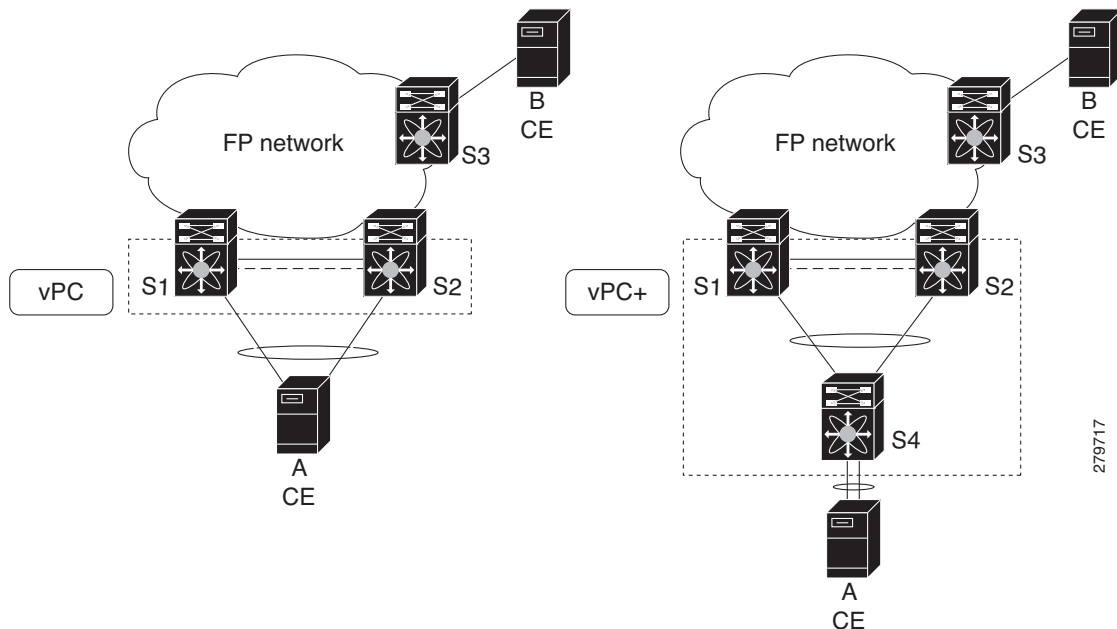
A vPC+ domain enables a Cisco Nexus device that has FabricPath enabled to form a single vPC+, which is a unique virtual switch to the rest of the FabricPath network. You configure the same domain on each switch to enable the peers to identify each other and to form the vPC+. Each vPC+ has its own virtual switch ID.

**Note**

We do not recommend that you enable the vPC peer switch feature when you are using a vPC+.

A vPC+ must still provide active-active Layer 2 paths for dual-homed CE switches or clouds, even though the FabricPath network allows only 1-to-1 mapping between the MAC address and the switch ID. vPC+ provides the solution by creating a unique virtual switch to the FabricPath network (see Figure 3-2).

Figure 3-2 Comparison of vPC and vPC+



The FabricPath switch ID for the virtual switch becomes the outer source MAC address (OSA) in the FabricPath encapsulation header. Each vPC+ domain must have its own virtual switch ID.

Layer 2 multipathing is achieved by emulating a single virtual switch. Packets forwarded from host A to host B are sent to the MAC address of the virtual switch as the transit source, and traffic from host B to host A is now load balanced.

The vPC+ downstream links are FabricPath edge interfaces that connect to the CE hosts.

The First Hop Routing Protocols (FHRPs), which include the Hot Standby Routing Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP), interoperate with a vPC+. You should dual-attach all Layer 3 switches to both vPC+ peer switches.

**Note**

You must enable the Layer 3 connectivity from each vPC+ peer switch by configuring a VLAN network interface for the same VLAN from both switches.

Both the primary and secondary vPC+ switches forward traffic, but only the primary FHRP switch responds to ARP requests.

To simplify initial configuration verification and vPC+/HSRP troubleshooting, you can configure the primary vPC+ peer switch with the FHRP active router highest priority.

In addition, you can use the `priority` command in the `if-hsrp` configuration mode to configure failover thresholds for when a group state enabled on a vPC+ peer is in standby or in listen state. You can configure lower and upper thresholds to prevent the group state flap, if there is an interface flap (this feature is useful when there is more than one tracking object per group).

When the primary vPC+ peer switch fails over to the secondary vPC+ peer switch, the FHRP traffic continues to flow seamlessly.

You should configure a separate Layer 3 link for routing from the vPC+ peer switches, rather than using a VLAN network interface for this purpose.

**Caution**

We do not recommend that you configure the burnt-in MAC address option (`use-bia`) for hot standby router protocol HSRP or manually configure virtual MAC addresses for any FHRP protocol in a vPC+ environment because these configurations can adversely affect the vPC+ load balancing.

The HSRP `use-bia` is not supported with a vPC+. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC+ peer switches.

You can configure a restore timer that delays the vPC+ coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature allows you to avoid packet drops if the routing tables do not converge before the vPC+ is once again passing traffic.

Use the **`delay restore`** command to configure this feature.

See the *Cisco Nexus 6000 Series NX-OS Unicast Routing Configuration Guide, Release 6.0* for more information on FHRPs and routing.

Anycast HSRP

Cisco NX-OS provides a way to facilitate further scalability at the spine layer giving support for more than two nodes. You can create an anycast bundle that is an association between a set of VLANs and an anycast switch ID. An anycast switch ID is the same as an emulated switch ID except the anycast switch ID is shared across more than two gateways. The set of VLANs or HSRP group elects an active router and a standby router. The remaining routers in the group are in listen state.

The active HSRP router advertises the anycast switch ID as the source switch ID in FabricPath IS-IS. The leaf switches learn that the anycast switch ID is reachable by all of the routers in the group.

Cisco NX-OS supports a maximum of four gateways. All the first-hop gateways at the spine layer must function in active-active forwarding mode. IP packets are received by any of the spine switches with the destination set as the gateway MAC address and these packets are terminated and locally forwarded.

Guidelines and Limitations for FabricPath Interfaces

FabricPath has the following configuration guidelines and limitations:

- An anycast HSRP bundle provides support for a maximum of four nodes at the spine layer.
- An anycast HSRP bundle is supported only in HSRP version 2.
- A maximum of 10 anycast bundles are supported.

Configuring FabricPath Interfaces

You must configure FabricPath interfaces on all switches that are enabled for FabricPath.



Tip

If you cannot see any of these commands, verify that you have installed and enabled the Fabricpath feature set.

This section includes the following topics:

- [Configuring FabricPath Interfaces, page 3-28](#)
- [Configuring the STP Priority with Rapid PVST+, page 3-30](#)
- [Configuring the STP Priority with MST, page 3-31](#)
- [Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch \(Optional\), page 3-33](#)
- [Configuring a vPC+ Switch ID, page 3-34](#)
- [Configuring an Anycast HSRP Bundle, page 3-35](#)

Configuring FabricPath Interfaces

You must configure the interfaces for the FabricPath network as FabricPath interfaces.

BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

SUMMARY STEPS

1. **configure terminal**
2. **interface** [*ethernet slot/port* | **port-channel** *channel-no*]
3. **fabricpath topology** *topology-number*
4. [**no**] **switchport mode fabricpath**
5. (Optional) **system default switchport fabricpath**
6. **exit**
7. (Optional) **show interface**
8. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|--|
| Step 1 | <pre>configure terminal</pre> <p>Example: switch# configure terminal switch(config)#</p> | Enters global configuration mode. |
| Step 2 | <pre>interface [ethernet slot/port port-channel channel-no]</pre> <p>Example: switch(config)# interface ethernet 2/11-15 switch(config-if)# </p> | <p>Enters interface configuration mode and specifies the interfaces that you want to configure as FabricPath.</p> <p>The slot can be from 1 to 3. The following list defines the slots available:</p> <ul style="list-style-type: none"> • Slot 1 includes all the fixed ports. A Fabric Extender only has one slot. • Slot 2 includes the ports on the upper expansion module (if populated). • Slot 3 includes the ports on the lower expansion module (if populated). <p>The port number within a particular slot can be from 1 to 128.</p> <p>The port channel number assigned to the EtherChannel logical interface can be from 1 to 4096.</p> <p>Note If this is a 10G breakout port, the <i>slot/port</i> syntax is <i>slot/QSFP-module/port</i>.</p> |
| Step 3 | <pre>fabricpath topology topology-number</pre> <p>Example: switch(config-if)# fabricpath topology 1 switch(config-if)# </p> | <p>Specifies the FabricPath topology number. The topology number can be from 1 to 63.</p> <p>Note Cisco Nexus Series only supports two topologies; the default or base topology (topology 0) and another optional topology (for example, topology 1).</p> |
| Step 4 | <pre>[no] switchport mode fabricpath</pre> <p>Example: switch(config-if)# switchport mode fabricpath switch(config-if)# </p> | <p>Specifies interfaces as FabricPath ports.</p> <p>Note The no keyword returns the interface to the default CE access interface. The FabricPath ports carry traffic only for those VLANs that are configured as FabricPath VLANs.</p> |
| Step 5 | <pre>system default switchport fabricpath</pre> <p>Example: switch(config-if)# system default switchport fabricpath switch(config-if)# </p> | (Optional) Converts all CE interfaces on the switch to FabricPath interfaces simultaneously. |
| Step 6 | <pre>exit</pre> <p>Example: switch(config-if)# exit switch(config)# </p> | Exits interface configuration mode. |

| | Command | Purpose |
|--------|--|---|
| Step 7 | show interface Example: switch(config)# show interface | (Optional) Displays information on all interfaces. |
| Step 8 | copy running-config startup-config Example: switch(config)# copy running-config startup-config | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure specified interfaces as FabricPath interfaces:

```
switch# configure terminal
switch(config)# interface ethernet 2/11-15
switch(config-if)# switchport mode fabricpath
switch(config-if)#
```

Configuring the STP Priority with Rapid PVST+

All Layer 2 gateway switches must have the same bridge priority when they are in the same STP domain. Make sure that the STP priority configured for the Layer 2 gateway switches on a FabricPath network is the lowest value in the Layer 2 network. Additionally, the priorities must match.

We recommend that you configure the STP priority on all FabricPath Layer 2 gateway switches to 8192.

BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree pseudo-information**
3. **spanning-tree vlan [vlan-id] root priority [value]**
4. **exit**
5. **(Optional) show spanning-tree summary**
6. **(Optional) copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>spanning-tree pseudo-information</code> Example: switch(config)# <code>spanning-tree pseudo-information</code> switch(config-pseudo)# | Enters spanning tree pseudo-information configuration mode. |
| Step 3 | <code>spanning-tree vlan [vlan-id] root priority [value]</code> Example: switch(config-pseudo)# <code>spanning-tree vlan 11-20 root priority 8192</code> | You must configure all the Rapid PVST+ VLANs on all the FabricPath Layer 2 gateway interfaces to a lower STP priority. We recommend that you configure the priority to be 8192. |
| Step 4 | <code>exit</code> Example: switch(config-pseudo)# <code>exit</code> switch(config)# | Exits the spanning-tree pseudo-information configuration mode. |
| Step 5 | <code>show spanning-tree summary</code> Example: switch(config)# <code>show spanning-tree summary</code> | (Optional) Displays STP information. |
| Step 6 | <code>copy running-config startup-config</code> Example: switch(config)# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure the Rapid PVST+ VLANs on the FabricPath Layer 2 gateway switches to have an STP priority of 8192:

```
switch# configure terminal
switch(config) spanning-tree pseudo-information
switch(config-pseudo)# spanning-tree vlan 11-20 root priority 8192
switch(config)#
```

See the *Cisco Nexus 6000 Series NX-OS Layer 2 Switching Configuration Guide, Release 6.0* for more information on this command.

Configuring the STP Priority with MST

All Layer 2 gateway switches must have the same bridge priority when they are in the same STP domain. Make sure that the STP priority configured for the Layer 2 gateway switches on a FabricPath network is the lowest value in the Layer 2 network. Additionally, the priorities must match.

We recommend that you configure the STP priority for all Multiple Spanning-Tree (MST) instances on all FabricPath Layer 2 gateway switches to 8192.

BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree pseudo-information**
3. **spanning-tree mst [instance-id] priority [value]**
4. **exit**
5. **(Optional) show spanning-tree summary**
6. **(Optional) copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters global configuration mode. |
| Step 2 | spanning-tree pseudo-information Example: switch(config)# spanning-tree pseudo-information switch(config-pseudo)# | Enters spanning tree pseudo-information configuration mode. |
| Step 3 | spanning-tree mst [instance-id] priority [value] Example: switch(config-pseudo)# spanning-tree mst 1-5 priority 8192 | Configures all the MST VLANs on all the FabricPath Layer 2 gateway interfaces to a lower STP priority. We recommend that you configure the priority to be 8192. |
| Step 4 | exit Example: switch(config-pseudo)# exit switch(config)# | Exits spanning tree pseudo-information configuration mode. |
| Step 5 | show spanning-tree summary Example: switch(config)# show spanning-tree summary | (Optional) Displays information on STP. |
| Step 6 | copy running-config startup-config Example: switch(config)# copy running-config startup-config | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure the MST instances on the FabricPath Layer 2 gateway switches to have an STP priority of 8192:

```
switch# configure terminal
```

```
switch(config)# spanning-tree pseudo-information
switch(config-pseudo)# spanning-tree mst 1-5 priority 8192
switch(config-pseudo)#
```

See the *Cisco Nexus 6000 Series NX-OS Layer 2 Switching Configuration Guide, Release 6.0* for more information on this command.

Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch (Optional)

Because there can be many FabricPath Layer 2 gateway switches attached to a single FabricPath network, there are also many separate STP domains that are each connected to a Layer 2 gateway switch. You can configure a unique STP domain ID in the FabricPath network to propagate TCNs across all the STP domains that are connected to the FabricPath network. This configuration ensures that all MAC addresses are flushed when the software receives a TCN.

BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree domain *domain-id***
3. **exit**
4. **(Optional) show spanning-tree summary**
5. **(Optional) copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|--|---|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters global configuration mode. |
| Step 2 | spanning-tree domain <i>domain-id</i> Example: switch(config)# spanning-tree domain 5 switch(config)# | Assigns an STP domain ID to the different STP domains attached to FabricPath Layer 2 gateway switches that are connected to a single FabricPath network. The range is from 1 to 1023. |
| Step 3 | exit Example: switch(config)# exit switch# | Exits configuration mode. |

| | Command | Purpose |
|--------|---|---|
| Step 4 | <code>show spanning-tree summary</code> Example: <code>switch# show spanning-tree summary</code> | (Optional) Displays STP information. |
| Step 5 | <code>copy running-config startup-config</code> Example: <code>switch# copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure the STP domain ID attached to the FabricPath Layer 2 gateway switch:

```
switch# configure terminal
switch(config)# spanning-tree domain 5
switch(config)# exit
switch
```

Configuring a vPC+ Switch ID

You can configure a vPC+ switch ID.



Note

The vPC+ virtual SID cannot be the same as the FabricPath SID for the switch.

BEFORE YOU BEGIN

Ensure that you have enabled the FabricPath feature.

SUMMARY STEPS

1. `configure terminal`
2. `vPC domain domain-id`
3. `fabricpath switch-id switch-id`

DETAILED STEPS

| | Command | Purpose |
|--------|---|-----------------------------------|
| Step 1 | <code>configure terminal</code> Example: <code>switch# configure terminal</code> | Enters global configuration mode. |

| | Command | Purpose |
|--------|---|---|
| Step 2 | vpc domain <i>domain-id</i> Example: switch(config)# vpc domain 1 | Creates a vPC+ domain on the switch, and enters the vPC-domain configuration mode for configuration purposes. |
| Step 3 | fabricpath switch-id <i>switch-id</i> Example: switch(config-vpc-domain)# fabricpath switch-id 1 | Assigns a static vPC+ ID to the vPC+ peer. The range is from 1 to 4094. This static ID is the virtual switch ID for FabricPath encapsulation. Note You must assign the same vPC+ switch ID to each of the two vPC+ peer switches before they can form an adjacency. |

This example shows how to configure a vPC+ switch ID on each vPC+ peer switch:

```
switch# configure terminal
switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath switch-id 1
```

See the *Cisco Nexus 6000 Series NX-OS Interfaces Configuration Guide, Release 6.0* for complete information on configuring vPCs.

Configuring an Anycast HSRP Bundle

Configuring an Anycast Bundle

You can create an anycast bundle that is an association between a set of VLANs and an anycast switchID.

BEFORE YOU BEGIN

- Ensure that you have enabled the FabricPath feature.
- Ensure that you have enabled the HSRP feature.
- Ensure that you have enabled the interface VLAN feature.

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch# configure terminal</code> | Enters global configuration mode. |
| Step 2 | <code>switch(config)# [no] hsrp anycast bundle-id {ipv4 ipv6 both}</code> | Configures an anycast bundle. The arguments and keywords are as follows: <ul style="list-style-type: none"> <i>bundle-id</i>—Bundle ID. The range is from 1 to 4096. ipv4—Specifies an IPv4 bundle. All the IPv4 groups in the interface are associated with this bundle. ipv6—Specifies an IPv6 bundle. All the IPv6 groups in the interface are associated with this bundle. both—Specifies an IPv4 and IPv6 bundle. This is the default. All the IPv4 and IPv6 groups in the interface are associated with this bundle. |
| Step 3 | <code>switch(config-anycast-bundle)# [no] force-gateway down</code> | Enforces the anycast bundle to remain in the down state even if one invalid VLAN is configured for the bundle. The bundle will come up only when all configured VLANs are valid. This mode favors loss of traffic. If you use the no form of the command, the bundle will come up even if one valid VLAN configuration exists. This mode favors loss of anycast gateway. |
| Step 4 | <code>switch(config-anycast)# [no] switch-id asid</code> | Configures the switch ID for the anycast bundle. <ul style="list-style-type: none"> <i>asid</i>—Specifies the anycast switch to be used for the bundle. |
| Step 5 | <code>switch(config-anycast)# [no] vlan range</code> | Configures the VLAN range for the anycast bundle. <ul style="list-style-type: none"> <i>range</i>—Specifies the VLAN range. |
| Step 6 | <code>switch(config-anycast)# [no] priority priority_value</code> | Configures the priority for the anycast bundle. The <i>priority_value</i> is used to elect a root for all the groups in the range. The range is from 1 to 127 and the default value is 100. |
| Step 7 | <code>switch(config-anycast)# [no] track object_id</code> | Configures the tracking value that is used to track the anycast bundle. The range of <i>object_id</i> is from 1 to 500. The default value is 0 which indicates that nothing is tracked. |
| Step 8 | <code>switch(config-anycast)# [no] timer hello_interval</code> | Configures the timers for the groups using this anycast bundle. The default value of <i>hello_interval</i> is 3. |
| Step 9 | <code>switch(config-anycast)# [no] shutdown</code> | Configures the group to take the switch out of the anycast bundle. To bring the switch into the anycast bundle, enter the no form of the command. |

| | Command | Purpose |
|---------|---|---|
| Step 10 | <code>switch(config-anycast)# show hsrp anycast [bundle_id ipv4 ipv6 both]</code> | (Optional) Displays anycast HSRP bundle information. |
| Step 11 | <code>switch(config-anycast)# copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to configure an anycast bundle for a selection of VLANs:

```
switch# configure terminal
switch(config)# hsrp anycast 1 ipv4
switch(config-anycast-bundle)# force gateway-down
switch(config-anycast-bundle)# switch-id 1300
switch(config-anycast-bundle)# vlan 1, 20-30
switch(config-anycast-bundle)# priority 90
switch(config-anycast-bundle)# track 2
switch(config-anycast-bundle)# timer 15
switch(config-anycast-bundle)# no shutdown
```

Verifying the FabricPath Interface Configuration

To display FabricPath interfaces information, perform one of the following tasks:

| Command | Purpose |
|---|--|
| <code>show feature-set</code> | Displays whether FabricPath is enabled on the switch or not. |
| <code>show interface brief</code> | Displays information on all interfaces. |
| <code>show interface switchport</code> | Displays information, including access and trunk interface, for all the Layer 2 interfaces. |
| <code>show interface type {slot/port \ channel-number} [trunk]</code> | Displays interface configuration information. Note If this is a 10G breakout port, the <i>slot/port</i> syntax is <i>slot/QSFP-module/port</i> . |
| <code>show interface capabilities</code> | Displays information on the capabilities of the interfaces. |
| <code>show interface status</code> | Displays information on the status of the interfaces. |
| <code>show spanning-tree summary</code> | Displays STP information. |
| <code>show fabricpath is-is database</code> | Displays STP TCN information. |
| <code>show vPC brief</code> | Displays brief information on the vPC+ domains. |
| <code>show vpc consistency paramters {[vpc vpc-id global vlans interface {[ethernet mod/slot port-channel channel number]}]}</code> | Displays the status of those parameters that must be consistent across all vPC+ domain interfaces. |
| <code>show vPC peer-keepalive</code> | Displays information on the peer-keepalive messages. |

| Command | Purpose |
|--|---|
| show vPC role | Displays the peer status, the role of the local switch, the vPC+ domain's MAC address and system priority, and the MAC address and priority for the local vPC+ domain's switch. |
| show vPC statistics {[peer-keepalive peer-link vpc vpc-id]} | Displays statistics on the vPC+ domains. |
| show running-config vPC [all] | Displays running configuration information for vPCs and vPC+ domains. |
| show hsrp anycast [bundle_id ipv4 ipv6][brief] | Displays information for anycast bundles. |
| show hsrp anycast brief | Displays information for anycast bundles. |
| show hsrp anycast interface vlan interface | Displays information about the interface in the anycast bundle. |
| show hsrp anycast summary | Displays the summary of anycast information. |
| show hsrp anycast internal info bundle_id ipv4 ipv6 | Displays all the data structures related to anycast. |
| show hsrp anycast remote-db [bundle_id ipv4 ipv6] | Displays the remote database for all the bundles. |

Monitoring FabricPath Interface Statistics

Use the following commands to display or clear FabricPath interface statistics:

- **show interface counters** [module module]
- **show interface counters detailed** [all]
- **show interface counters errors** [module module]
- **clear counters** [interface {[all | ethernet mod/slot | port-channel channel-number | loopback virtual-interface-num | mgmt mgmt-interface-num | vlan vlan-id]}]

Configuration Example for FabricPath Interface

To configure FabricPath interfaces, perform the following tasks on each switch:

- Enable FabricPath on each switch.
- Configure the interfaces that you want to designate as FabricPath interfaces.
- Set the STP priority switch to 8192 on all FabricPath Layer 2 gateway switches.
- (Optional) Set the STP domain ID for each of the separate STP domains that are connected to the FabricPath network.
- (Optional) Configure a vPC+ switch ID.

To configure FabricPath interfaces, follow these steps:

Step 1 (Optional) Enable FabricPath on each switch.

```
switch# configure terminal
```



```
switch(config)# feature-set fabricpath
switch(config-lldp)# exit
switch(config)#
```

Step 2 After you enable FabricPath on the switch, configure the specified interface as FabricPath interfaces.

```
switch(config)# interface ethernet 1/2
switch(config-if)# switchport mode fabricpath
switch(config-if)# exit
switch(config)#
```

Step 3 Configure the STP priority for all Rapid PVST+ VLANs as 8192.

```
switch# configure terminal
switch(config)# spanning-tree pseudo-information
switch(config-pseudo)# spanning-tree vlan 11-20 priority 8192
switch(config-pseudo)#
```

Step 4 Configure the STP priority for all MST instances as 8192.

```
switch# configure terminal
switch(config)# spanning-tree pseudo-information
switch(config-pseudo)# spanning-tree mst 1-5 priority 8192
switch(config-pseudo)#
```

Step 5 (Optional) Configure the STP domain ID on each FabricPath Layer 2 gateway switch attached to the FabricPath network.

```
switch# configure terminal
switch(config)# spanning-tree domain 5
switch(config)
```

Step 6 (Optional) Configure the vPC+ switch ID.

```
switch# configure terminal
switch(config)# vPC domain 5
switch(config-vPC-domain)# fabricpath switch-id 100
switch(config-vPC-domain)# exit
switch(config)
```

If you are configuring the vPC+ with no existing vPC+, follow these steps:

1. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.
2. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.
3. On the vPC+ peer link port channel, enter the **vPC peer-link** command.

If you are changing an existing vPC configuration to a vPC+, follow these steps:

1. On each vPC peer link port channel, enter the **shutdown** command.
2. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.
3. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.
4. On the vPC+ peer link port channel, enter the **no shutdown** command.

Step 7 Save the configuration.

```
switch(config)# copy running-config startup-config
switch(config)#
```

If you are configuring a vPC+, and you see the following situations, you must enter the **shutdown** *command* and then the **no shutdown** *command* on all the peer-link interfaces:

- There is no switchport mode FabricPath configuration on the peer-link interfaces, but the FabricPath switch ID is configured in the vPC domain.
- The **switchport mode fabricpath** configuration is on the peer-link interfaces, but there is no FabricPath switch ID in the vPC domain.

Feature History for Configuring FabricPath Interfaces

Table 3-1 lists the release history for this feature.

Table 3-1 Feature History for FabricPath

| Feature Name | Releases | Feature Information |
|-------------------------------|-------------|---|
| Anycast HSRP | 7.1(0)N1(1) | Added the ability to add or delete a VLAN to or from an existing VLAN range (for an HSRP Anycast bundle) without having to enter the complete VLAN range again. |
| Anycast HSRP and overload bit | 7.1(0)N1(1) | The anycast switch ID is no longer advertised when the FabricPath Layer 2 IS-IS overload-bit is set. Please see more details about the Fabricpath Layer IS-IS overload bit in the section "Configuring Advanced FabricPath Features." |
| FabricPath Interfaces | 6.0(2)N1(1) | This feature was introduced. |



Configuring FabricPath Forwarding

This chapter contains the following sections:

- [Information About FabricPath Forwarding, page 4-41](#)
- [Configuring FabricPath Forwarding, page 4-45](#)
- [Verifying the FabricPath Configuration, page 4-49](#)
- [Configuration Example for FabricPath Forwarding, page 4-50](#)
- [Default Settings for Fabric Path Forwarding, page 4-51](#)
- [Feature History for Configuring FabricPath, page 4-51](#)



Note

For information about prerequisites, guidelines and limitations, and licensing requirements for FabricPath, see [Chapter 1, “Overview.”](#)

Information About FabricPath Forwarding

This section includes the following sections:

- [FabricPath Forwarding Overview, page 4-41](#)
- [FabricPath VLANs, page 4-42](#)
- [Forwarding Known Unicast Packets Using ECMP, page 4-43](#)
- [Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets, page 4-44](#)

FabricPath Forwarding Overview

FabricPath provides a multipath Layer 2 domain that does not require STP for a loop-free environment. Using the Intermediate System-to-Intermediate System (IS-IS) protocol, the switch provides multiple paths for Layer 2 packets.

Each FabricPath interface can learn multiple parallel paths to the other nodes in the FabricPath network. Because you do not need to use STP, all the paths are available for forwarding traffic. The switch assigns the optimal path per flow.

The flow for known unicast packets is determined by the hierarchical FabricPath outer destination address (ODA) and the outer source address (OSA) value (see [Chapter 2, “Configuring FabricPath Switching,”](#) for more information on FabricPath hierarchical encapsulation). The software uses IS-IS equal cost multipath (ECMP) to choose the forwarding path for these flows.

For multidestination traffic (unknown unicast, broadcast, and multicast), the software creates two paths or trees. The broadcast and unknown unicast traffic flows through one of these trees. The software distributes the multicast traffic between the two trees based on a hash. The software load balances multicast traffic in the FabricPath network (see the [“Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets”](#) section on page 4-44 for more information).

FabricPath Layer 2 IS-IS defines the trees. The highest switch ID is chosen for the root and the first tree is built with that switch as the root. (If needed, you can configure the switch ID to specify which switch has the highest switch ID and is, therefore, the root for the first tree.) Then, the root switch for the first tree elects the root of the second tree, again based on switch ID, and the second tree flows from that root switch. All of this information is advertised to the FabricPath network using Layer 2 IS-IS, so all the switches in the network have the same information.

The software assigns the path at ingress and encodes that path in the FTag portion of the FabricPath header. The software assigns one FTag per tree. Once decided and tagged, the packet uses the same tree throughout the entire FabricPath network. All the nodes in the FabricPath network forward traffic based on this same information because all nodes have the same information using Layer 2 IS-IS.

The FabricPath frame has a Reverse Path Forwarding (RPF) mechanism for multidestination packets, which verifies that the packet is arriving on an interface that leads to the source switch. RPF drops the packet if it is received from an interface that is not part of the tree.

The FabricPath Layer 2 IS-IS protocol floods the link-state information across the FabricPath network. Each switch sends hello packets on each FabricPath link and discovers its neighbors. When a neighbor is discovered, the software creates an IS-IS adjacency. Each switch also sends advertisements and updates to the link-state database through all the existing adjacencies.


FabricPath VLANs

To interact with the Classical Ethernet (CE) network, you set VLANs to either CE or FabricPath mode. The CE VLANs carry traffic from the CE hosts to the FabricPath interfaces, and the FabricPath VLANs carry traffic throughout the FabricPath topology. Only the active FabricPath VLANs configured on a switch are advertised as part of the topology in the Layer 2 Intermediate System-to-Intermediate System (IS-IS) messages.

The software automatically assigns all FabricPath interfaces and FabricPath VLANs to the FabricPath topology. No added configuration is required. All the FabricPath VLANs and interfaces belong to that topology.

Figure 4-1 shows an example of the FabricPath topology with Classical Ethernet switches and FabricPath/CE VLANs.

Figure 4-1 Example FabricPath Topology and Classical Ethernet Hosts



The default VLAN mode on the switch is the CE VLAN mode. The FabricPath interfaces carry traffic only on the FabricPath VLANs; the CE VLANs do not come up on these interfaces.

You must exit the VLAN configuration mode for the VLAN mode change to take effect.



Note

After you configure the VLANs and interfaces, no further configuration is required. The software automatically creates and assigns the paths, and provides load balancing.

Forwarding Known Unicast Packets Using ECMP

The switch forwards unicast traffic per flow using the ODA field in the FabricPath header for known unicast traffic. The FabricPath-enabled switch assigns the switch ID and the ODA for all encapsulated traffic at the ingress switch. (See Chapter 2, “Configuring FabricPath Switching,” for more information on FabricPath encapsulation.)

Once the switch assigns the ODA, the FabricPath switch uses the FabricPath Layer 2 IS-IS ECMP to forward known unicast traffic. FabricPath Layer 2 IS-IS has up to 16 active Layer 2 paths for all known unicast packets. The Layer 2 IS-IS messages used by FabricPath are separate and distinct from the Layer 3 IS-IS messages used by the routing protocols and the Overlay Transport Virtualization (OTV).

The switches within the FabricPath network exchange topology information using IS-IS adjacencies and forward the traffic along those paths for known unicast traffic flows. Each node in the FabricPath network looks at the FabricPath header for each traffic flow and makes an ECMP forwarding choice based on the available next hops.

Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets

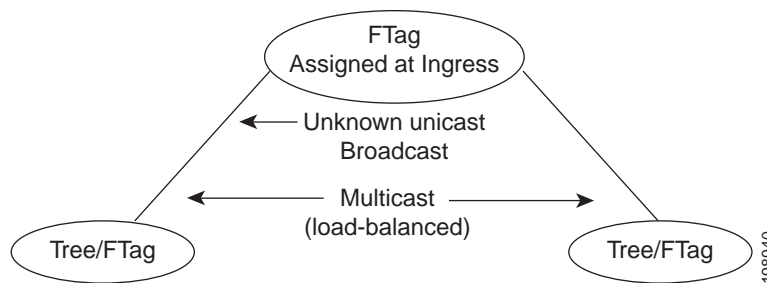
FabricPath has a loop-free broadcast functionality that carries broadcast, unknown unicast, and multicast packets, or multidestination traffic. For each broadcast, unknown unicast, and multicast traffic flow, the software chooses the forwarding path from among two software-created paths or trees. The software creates two trees to forward multidestination traffic.

For the FabricPath network, the software creates a broadcast tree that carries broadcast traffic, unknown unicast traffic through the FabricPath network. The software also creates another tree; all the multicast traffic flows are load balanced across these two trees for each flow. Each tree is identified in the FabricPath network by a unique value or FTag. Within the FabricPath network, the software elects a root node that becomes root for the broadcast tree. That node also identifies another bridge to become root for the second multidestination tree, which load balances the multicast traffic.

The FTag is assigned by the ingress switch, along with the ODA and OSA, as part of the FabricPath encapsulation. The FTag determines which loop-free tree the multidestination traffic flow follows through the FabricPath network. The software assigns the trees per flow.

Figure 4-2 shows these trees.

Figure 4-2 *Trees for Forwarding Multidestination FabricPath Flows for a Given Flow*



Each node in the FabricPath network shares the same view of the forwarding trees for a given FTag.

Forwarding Multicast Packets

Using FabricPath, you can configure Layer 2 multicast multipathing. FabricPath uses a hash-based system to assign each of the multicast flows to one of the two designated trees to ensure that the multicast traffic is load balanced.

The software uses FabricPath Layer 2 IS-IS and Classical Ethernet IGMP snooping to learn the multicast group information at the boundaries of the FabricPath/Classical Ethernet network. The software carries that information through the FabricPath network using a new Layer 2 IS-IS LSP called Group Membership LSP (GM-LSP). GM-LSPs carry multicast group/source membership information. This information is carried across the FabricPath network. All FabricPath switches maintain multicast routing information and forward multicast data packets only to switches that have interested receivers. Each node in the FabricPath topology shares the same view and contains the same information.

The multicast traffic uses VLAN information to allocate traffic to one or the other of the two multicast trees. This software constrains multicast based on the group IP address

IGMP snooping and FabricPath IS-IS, using GM-LSP, work together to build per-VLAN multicast group-based trees across the FabricPath network. IGMP snooping on edge switches learns of receivers and routers and builds an edge-port multicast state. FabricPath Layer 2 IS-IS propagates this group

information through the FabricPath network using GM-LSPs and build the state in the FabricPath network. These GM-LSPs originate from the switches at the edge of the FabricPath network that have multicast groups.

For Layer 2 multicast traffic, you need not run PIM at all when using FabricPath.

For Layer 3 multicast packets, the software sets the ODA to a special multicast group that identifies all IP routers for that group and forwards the traffic along the tree for that group.

Configuring FabricPath Forwarding



Tip

If you cannot see any of these commands, verify that you have installed and enabled the FabricPath feature set on the switch.

You must make these configurations on every switch in the FabricPath network.

Only those VLANs that are configured as FabricPath VLANs can belong to the FabricPath topology. By default, all FabricPath VLANs are assigned to the base topology (Topology 0). A FabricPath VLAN can be configured to be assigned to a different topology.

FabricPath has two topologies, you need to set the VLAN mode for those VLANs that you want to traverse the FabricPath network to FabricPath VLAN.

After you configure the VLAN modes and interfaces, the software automatically creates the required paths. You do not need to perform any additional configuration for FabricPath.

This section includes the following topics:

- [Setting the VLAN Mode to FabricPath or CE, page 4-45](#)
- [Mapping FabricPath VLANs to a Topology, page 4-46](#)

Setting the VLAN Mode to FabricPath or CE

The default VLAN mode is CE.

You designate those VLANs that you want to carry FabricPath traffic on the network by configuring them as FabricPath VLANs.

BEFORE YOU BEGIN

Ensure that you have created the VLANs.

SUMMARY STEPS

1. **configure terminal**
2. **vlan *vlan-id***
3. **mode [ce | fabricpath]**
4. **exit**
5. **(Optional) show fabricpath topology vlans [active]**
6. **(Optional) copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)# | Enters global configuration mode. |
| Step 2 | <code>vlan vlan-id</code> Example: switch(config)# <code>vlan 1-10</code> switch(config-vlan)# | Enters the VLAN configuration mode and identifies those VLANs that you want to carry FabricPath traffic. The VLAN range is from 1 to 4094. Note You cannot modify VLAN 1 or any of the internally allocated VLANs. |
| Step 3 | <code>mode [ce fabricpath]</code> Example: switch(config-vlan)# <code>mode fabricpath</code> switch(config-vlan)# | Configures the VLANs as FabricPath VLANs. The default VLAN mode is CE. Note A VLAN must be either a CE or an FabricPath VLAN on the FabricPath switch. |
| Step 4 | <code>exit</code> Example: switch(config-vlan)# <code>exit</code> switch(config)# | Exits the VLAN configuration mode. Note As with all VLANs, you must exit the VLAN configuration mode for the VLAN mode (CE or FabricPath) to take effect. |
| Step 5 | <code>show fabricpath topology vlans [active]</code> Example: switch# <code>show fabricpath topology vlans</code> | (Optional) Displays information about all active VLANs in the FabricPath topology. |
| Step 6 | <code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code> | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to specify VLANs as FabricPath VLANs:

```
switch# configure terminal
switch(config)# vlan 1-10
switch(config-vlan)# mode fabricpath
switch(config-vlan)# exit
switch(config)# exit
```

Mapping FabricPath VLANs to a Topology

The default VLAN mode is CE.

You designate those VLANs that you want to carry FabricPath traffic on the network by configuring them as FabricPath VLANs. By default, all FabricPath VLANs belong to default topology 0. Also FabricPath VLANs can be a member of only one topology.

BEFORE YOU BEGIN

Ensure that you have created the VLANs.

SUMMARY STEPS

1. **configure terminal**
2. **fabricpath topology *topology-number***
3. **member vlan *vlan-id***
4. **(Optional) show fabricpath topology vlans [active]**
5. **(Optional) copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|--|---|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters global configuration mode. |
| Step 2 | fabricpath topology <i>topology-number</i> Example: switch(config)# fabricpath topology 1 switch(config-fp-topology)# | Specifies the FabricPath topology number. The topology number can be from 1 to 63. Note Cisco Nexus 6000 Series switches support two Fabricpath topologies, the default or base topology (topology 0), and another optional topology (for example, topology 1). |
| Step 3 | member vlan <i>vlan-id</i> Example: switch(config-fp-topology)# member vlan 501-600 switch(config-vlan)# | Specifies the VLANs to assign to FabricPath topology. |
| Step 4 | show fabricpath topology vlans [active] Example: switch# show fabricpath topology vlans | (Optional) Displays information about all active VLANs in the FabricPath topology. |
| Step 5 | copy running-config startup-config Example: switch# copy running-config startup-config | (Optional) Copies the running configuration to the startup configuration. |

This example shows how to map a VLAN to a FabricPath topology:

```
switch# configure terminal
switch(config)# fabricpath topology 1
switch(config-fp-topology)# member vlan 501-600
switch(config-fp-topology)# exit
```

Configuring FabricPath Unicast Load Balancing (Optional)

The FabricPath network automatically balances unicast traffic when multiple paths are available. However, you can configure specific load balancing for the unicast traffic. The default is to use all options.

BEFORE YOU BEGIN

Ensure that you have enabled the FabricPath feature set.

SUMMARY STEPS

1. **configure terminal**
2. **fabricpath load-balance unicast {include-vlan | layer2 | layer3 | layer4 | mixed}**
3. **exit**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|---|
| Step 1 | <pre>configure terminal</pre> <p>Example: <pre>switch# configure terminal switch(config)#</pre></p> | Enter global configuration mode. |
| Step 2 | <pre>fabricpath load-balance unicast {include-vlan layer2 layer3 layer4 mixed}</pre> <p>Example: <pre>switch(config)# fabricpath load-balance unicast include-vlan switch(config-vlan)#</pre></p> | <p>Configures load balancing for FabricPath unicast traffic along the following parameters:</p> <ul style="list-style-type: none"> • include-vlan—Specifies that the load-balancing parameters use VLAN. • layer2—Specifies that the load-balancing parameters only include Layer 2 parameters. • layer3—Specifies that the load-balancing parameters only include Layer 3 parameters. • layer4—Specifies that the load-balancing parameters only include Layer 4 parameters. • mixed—Specifies that the load-balancing parameters include a mix of Layer 3 and Layer 4 parameters. This is the default configuration. <p>Note To return to the default unicast load-balancing scheme, enter the no form of this command.</p> |

| | Command | Purpose |
|--------|---|---|
| Step 3 | <code>exit</code> Example: <code>switch(config)# exit</code> <code>switch(config)#</code> | Exit configuration mode. |
| Step 4 | <code>copy running-config startup-config</code> Example: <code>switch# copy running-config startup-config</code> | (Optional) Copy the running configuration to the startup configuration. |

This example shows how to configure FabricPath unicast load balancing with source parameters:

```
switch# configure terminal
switch(config)# fabricpath load-balance unicast include-vlan
switch(config)#
```

Verifying the FabricPath Configuration

To display FabricPath forwarding information, perform one of the following tasks:



Tip

For more information about these commands and the options available, see the command reference for your platform.

| Command | Purpose |
|--|---|
| <code>show feature-set</code> | Displays whether FabricPath is enabled or not. |
| <code>show {I2 fabricpath} route [switchid <i>switch-id</i>] [detail] [hex]</code> | Displays unicast routes. |
| <code>show {I2 fabricpath} mroute {[vdc_omf] [vlan <i>vlanid</i>] {[omf] [flood] [source {<i>srcaddr</i> <i>ipv6srcaddr</i>}] [group {<i>groupaddr</i> <i>ipv6groupaddr</i>}]}} [resolved] [ftag <i>ftag-id</i>] [hex]}</code> | Displays multicast routes. |
| <code>show fabricpath topology [detail]</code> | Displays information on the FabricPath topology. |
| <code>show fabricpath topology interface</code> | Displays information on all FabricPath topology interfaces. |
| <code>show fabricpath topology vlan [active]</code> | Displays information on all FabricPath topology VLANs. |
| <code>show fabricpath topology ftag [active] [multicast] [unicast]</code> | Displays information on all FabricPath topology FTags. |
| <code>show running-config fabricpath</code> | Displays the running configuration for FabricPath. |

| Command | Purpose |
|--|--|
| show fabricpath load-balance unicast forwarding-path ftag <i>f-tag-ID</i> switchid <i>switch-ID</i> src-mac <i>MAC-addr</i> | Displays FabricPath unicast load-balancing information. This command contains the following options: <ul style="list-style-type: none"> <i>f-tag-ID</i>—FTag for which you want to view load-balancing. The FTag value is from 0 to 4,294,967,295. <i>switch-ID</i>—The ID of the switch for which you want to view load-balancing. The switch ID is from 0 to 4,294,967,295. |
| show fabricpath load-balance multicast ftag-selected vlan <i>vlan-id</i> macg <i>macg</i> | Displays FabricPath multicast load-balancing information. The output from this command is valid for multicast traffic ingressing on non-vPC ports. For multicast traffic ingressing on vPC ports, traffic picks the active Ftag only for that switch. This command contains the following options: <ul style="list-style-type: none"> <i>vlan-id</i>—A VLAN ID in a range between 1 and 4094. <i>macg</i>—The multicast group MAC address. |
| show vlan | Displays information on all FabricPath and CE VLANs. |

Configuration Example for FabricPath Forwarding

To configure the basic FabricPath network, you must accomplish the following tasks on each switch after you have configured the FabricPath interfaces:

- Enable the FabricPath feature set on each switch.
- Configure the FabricPath interfaces. (See [Chapter 3, “Configuring FabricPath Interfaces,”](#) for information on configuring FabricPath interfaces.)
- Configure the FabricPath VLANs. The default is CE VLANs.
- Enter the **show running-config fabricpath** command to make sure that your FabricPath configuration is correct.

To configure the FabricPath topology, follow these steps:

Step 1 Enable the FabricPath feature set.

```
switch# configure terminal
switch(config)# feature-set fabricpath
switch(config)#
```

Step 2 Set the VLAN modes for those VLANs that you want in the FabricPath topology to FabricPath.

```
switch# configure terminal
switch(config)# vlan 11-20
switch(config-vlan)# mode fabricpath
switch(config-vlan)# exit
```

```
switch(config)
```

Step 3 Display the configuration to ensure that you have the correct configuration.

```
switch(config)# show running-config fabricpath
switch(config)#
```

Step 4 Save the configuration.

```
switch(config)# copy running-config startup-config
switch(config)#
```

Default Settings for Fabric Path Forwarding

Table 4-1 lists the default settings for the FabricPath forwarding parameters

Table 4-1 Default FabricPath Parameters

| Parameters | Default |
|---------------------|------------|
| FabricPath Topology | 0 |
| VLAN mode | FabricPath |

Feature History for Configuring FabricPath

Table 4-2 lists the release history for this feature.

Table 4-2 Feature History for FabricPath

| Feature Name | Releases | Feature Information |
|-----------------------|-------------|-----------------------------|
| FabricPath Forwarding | 6.0(2)N1(1) | This feature was introduced |



Configuring Advanced FabricPath Features

This chapter contains the following sections:

- [Information About Advanced FabricPath Layer 2 IS-IS Configurations, page 5-53](#)
- [Setting Advanced FabricPath Layer 2 IS-IS Parameters, page 5-54](#)
- [Multi-Destination Tree, page 5-65](#)
- [Verifying the FabricPath Advanced Configurations, page 5-67](#)
- [Feature History for Configuring FabricPath Advanced Features, page 5-68](#)



Note

For information about prerequisites, guidelines and limitations, and licensing requirements for FabricPath, see [Chapter 1, “Overview.”](#)

Information About Advanced FabricPath Layer 2 IS-IS Configurations

We recommend that you run the FabricPath network using the default Layer 2 IS-IS configurations.

However, you can change many of the IS-IS settings as follows:

- Globally on the entire switch and on each switch in the FabricPath network
- On specified FabricPath switches within the FabricPath network

If you change any of the FabricPath Layer 2 IS-IS settings, ensure that you make the same changes for those global parameters on every switch in the FabricPath network and for those parameters on every applicable FabricPath switch in the network.

Layer 2 IS-IS is based on Layer 3 IS-IS with enhancements to run on Layer 2. The commands for Layer 2 IS-IS and Layer 3 IS-IS are not the same. Layer 2 IS-IS is the control plane in FabricPath and a single protocol controls all unicast and multicast traffic. From a forwarding standpoint, FabricPath Layer 2 IS-IS forwards traffic for unicast, unknown unicast, broadcast, and multicast frames. Using Layer 2 IS-IS, the software maintains loop-free paths throughout the FabricPath network. (see [Chapter 2, “Configuring FabricPath Switching,”](#) for information on default FabricPath Layer 2 IS-IS behavior and

Chapter 4, “Configuring FabricPath Forwarding,” for information on FabricPath forwarding.)

You can use these advanced FabricPath Layer 2 IS-IS configurations to fine-tune the operation of the FabricPath network.

Overload Bit

IS-IS uses the overload bit to tell other routers not to use the local router to forward traffic but to continue routing traffic destined for that local router.

You may want to use the overload bit in these situations:

- The router is in a critical condition.
- Graceful introduction and removal of the router to/from the network.
- Other (administrative or traffic engineering) reasons such as waiting for BGP convergence.

The Overload Bit feature is not backward compatible. For this feature to work properly, all FabricPath switches must support this feature.

You can configure the router to signal other routers not to use this router as an intermediate hop in their shortest path first (SPF) calculations. You can optionally configure the overload bit temporarily on startup, until BGP converges.

In addition to setting the overload bit, you might also want to suppress certain types of IP prefix advertisements from LSPs for Level 1 or Level 2 traffic.

Setting Advanced FabricPath Layer 2 IS-IS Parameters

Although the Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure parameters. Some FabricPath Layer 2 IS-IS parameters you configure globally and some you configure per switch. This section includes the following topics:

- [Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally, page 5-54](#)
- [Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface, page 5-59](#)
- [Clearing Advanced FabricPath Layer 2 IS-IS Counters, page 5-63](#)

Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally

Although the FabricPath Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure the global parameters.



Note

Graceful restart of IS-IS is not supported. While the **graceful-restart** command exists in the CLI, it is not supported at this time.

SUMMARY STEPS

1. **configure terminal**
2. **fabricpath domain default**
3. (Optional) **authentication-check**

4. (Optional) **authentication key-chain** *auth-key-chain-name*
5. (Optional) **authentication type** {cleartext | md5}
6. (Optional) **log-adjacency-changes**
7. (Optional) **lsp-gen-interval** *msecs* [*msecs msecs*]
8. (Optional) **lsp-mtu** *mtu*
9. (Optional) **max-lsp-lifetime** *secs*
10. (Optional) **maximum-paths** *max-paths*
11. (Optional) **reference-bandwidth** {*ref-mbps* [Mbps] | *ref-gbps* [Gbps]}
12. (Optional) **spf-interval** *msecs* [*msecs msecs*]
13. (Optional) **topology** *topology_number*
14. (Optional) **hostname** *dynamic*
15. (Optional) **root-priority** *value*
16. (Optional) **set-overload-bit** {always | on-start-up {*seconds* } }
17. **exit**
18. **exit**
19. (Optional) **show running-config**
20. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|--|--|
| Step 1 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 2 | fabricpath domain default Example: switch(config)# fabricpath domain default switch(config-fabricpath-isis)# | Enters the global FabricPath Layer 2 IS-IS configuration mode. |
| Step 3 | authentication-check Example: switch(config-fabricpath-isis)# authentication-check switch(config-fabricpath-isis)# | (Optional) Configures an authentication check when the switch receives a protocol data unit (PDU). To turn the authentication check off, enter the no form of this command. Note The default is ON. |

| | Command | Purpose |
|--------|--|---|
| Step 4 | <pre>authentication key-chain auth-key-chain-name</pre> <p>Example:</p> <pre>switch(config-fabricpath-isis)# authentication key-chain trees switch(config-fabricpath-isis)#</pre> | <p>(Optional) Configures the authentication keychain. To clear this parameter, enter the no form of this command.</p> <p>The maximum size for the name of an authentication keychain is 63 alphanumeric characters.</p> <p>An example of key chain creation is as follows:</p> <pre>key chain trees key 0 key-string cisco01 accept-lifetime 07:00:00 Sep 20 2011 infinite send-lifetime 07:00:00 Sep 20 2011 infinite</pre> <p>See the <i>Cisco Nexus 6000 Series NX-OS Security Configuration Guide, Release 6.0</i> for information about key chains.</p> |
| Step 5 | <pre>authentication-type {cleartext md5}</pre> <p>Example:</p> <pre>switch(config-fabricpath-isis)# authentication-type md5 switch(config-fabricpath-isis)#</pre> | <p>(Optional) Configures the authentication type. To clear this parameter, enter the no form of this command.</p> <p>You can set one of the following authentication types:</p> <ul style="list-style-type: none"> • cleartext—Specifies the cleartext authentication method. • md5—Specifies the Message Digest (MD5) authentication. |
| Step 6 | <pre>log-adjacency-changes</pre> <p>Example:</p> <pre>switch(config-fabricpath-isis)# log- adjacency-changes switch(config-fabricpath-isis)#</pre> | <p>(Optional) Sets the switch to send a log message when the state of a FabricPath Layer 2 IS-IS neighbor changes. To stop the log messages, enter the no form of this command. The default is OFF.</p> |

| | Command | Purpose |
|---------|--|--|
| Step 7 | <pre>lsp-gen-interval msec [msec msec]</pre> <p>Example: switch(config-fabricpath-isis)# lsp-gen-interval 100 switch(config-fabricpath-isis)#</p> | <p>(Optional) Configures the link-state packet (LSP) generation interval. To return to the default values, enter the no form of this command. The optional parameters are as follows:</p> <ul style="list-style-type: none"> • max-wait—The initial wait between the trigger and LSP generation. The range is from 50 to 120000 milliseconds, and the default value is 8000 milliseconds. • lsp-initial-wait—The initial wait between the trigger and LSP generation. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds. • lsp-second-wait—The second wait used for LSP throttle during backoff. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds. |
| Step 8 | <pre>lsp-mtu mtu</pre> <p>Example: switch(config-fabricpath-isis)# lsp-mtu 2000 switch(config-fabricpath-isis)#</p> | <p>(Optional) Sets the LSP MTU. To return to the default values, enter the no form of this command. The range is from 128 to 4352, and the default value is 1200.</p> |
| Step 9 | <pre>max-lsp-lifetime secs</pre> <p>Example: switch(config-fabricpath-isis)# max-lsp-lifetime 1000 switch(config-fabricpath-isis)#</p> | <p>(Optional) Sets the maximum LSP lifetime in seconds. To return to the default values, enter the no form of this command. The range is from 128 to 4352, and the default value is 1492.</p> |
| Step 10 | <pre>maximum-paths max-paths</pre> <p>Example: switch(config-fabricpath-isis)# maximum-paths 4 switch(config-fabricpath-isis)#</p> | <p>(Optional) Sets the maximum number of paths per destination. To return to the default values, enter the no form of this command. The range is from 1 to 16, and the default value is 16.</p> |
| Step 11 | <pre>reference-bandwidth {ref-mbps [Mbps] ref-gbps [Gbps]}</pre> <p>Example: switch(config-fabricpath-isis)# reference-bandwidth 200000 switch(config-fabricpath-isis)#</p> | <p>(Optional) Configures the reference bandwidth, which is used to assign the FabricPath Layer 2 IS-IS cost. The default value is 400000 Mbps. To return to the default values, enter the no form of this command. The optional parameters are as follows:</p> <ul style="list-style-type: none"> • ref-mbps—The range is from 1 to 4000000, and the default value is 4000000. • ref-gbps—The range is from 1 to 4000, and the default value is 400. |

| | Command | Purpose |
|---------|--|--|
| Step 12 | <p>spf-interval <i>msecs</i> [<i>msecs msecs</i>]</p> <p>Example: switch(config-fabricpath-isis)# spf-interval 10000 switch(config-fabricpath-isis)#</p> | <p>(Optional) Configures the interval between link-state advertisement (LSA) arrivals. To return to the default values, enter the no form of this command. The optional parameters are as follows:</p> <ul style="list-style-type: none"> • spf-max-wait—The maximum wait between the trigger and shortest path first (SPF) computation. The range is from 50 to 120000 milliseconds, and the default value is 8000 milliseconds. • spf-initial-wait—The initial wait between the trigger and SPF computation. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds. • spf-second-wait—The second wait used for SPF computation during backoff. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds. |
| Step 13 | <p>topology <i>topology_number</i></p> <p>Example: switch(config-fabricpath-isis)# topology 1 switch(config-fabricpath-isis)#</p> | <p>(Optional) Configures the topology number. The topology number can be from 1 to 63.</p> |
| Step 14 | <p>hostname <i>dynamic</i></p> <p>Example: switch(config-fabricpath-isis)# hostname dynamic switch(config-fabricpath-isis)#</p> | <p>(Optional) Enables dynamic hostname for the FabricPath Layer 2 IS-IS protocol. To disable the dynamic hostname, enter the no form of this command.</p> |
| Step 15 | <p>root-priority <i>value</i></p> <p>Example: switch(config-fabricpath-isis)# root-priority 100 switch(config-fabricpath-isis)#</p> | <p>(Optional) Sets the priority for which node becomes the Layer 2 IS-IS protocol root in the FabricPath network. The highest numerical value for the priority is likely to become the root. To return to the default values, enter the no form of this command. The range is from 1 to 255, and the default value is 64.</p> |
| Step 16 | <p>set-overload-bit {<i>always</i> <i>on-startup</i> {<i>seconds</i>}}</p> <p>Example: switch(config-fabricpath-isis)# set-overload-bit on-startup 30</p> | <p>(Optional) Sets the overload bit for IS-IS. The <i>seconds</i> range is from 5 to 86400.</p> |
| Step 17 | <p>exit</p> <p>Example: switch(config-fabricpath-isis)# exit switch(config)#</p> | <p>Exits global FabricPath Layer 2 IS-IS configuration mode.</p> |
| Step 18 | <p>exit</p> <p>Example: switch(config)# exit switch#</p> | <p>Exits global configuration mode.</p> |

| | Command | Purpose |
|---------|--|---|
| Step 19 | show running-config Example: switch# show running-config switch# | (Optional) Displays the running configuration. |
| Step 20 | copy running-config startup-config Example: switch# copy running-config startup-config | (Optional) Copies the running configuration to the startup configuration. |

See the *Cisco Nexus 6000 Series NX-OS Unicast Routing Configuration Guide, Release 6.0* for more information on IS-IS commands.

Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface

Although the FabricPath Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure the interface parameters.

SUMMARY STEPS

1. **configure terminal**
2. **interface** {*ethernet mod/slot* | *port-channel channel-number*}
3. (Optional) **fabricpath isis authentication-check**
4. (Optional) **fabricpath isis authentication key-chain** *auth-key-chain-name*
5. (Optional) **fabricpath isis authentication type** {*cleartext* | *md5*}
6. (Optional) **fabricpath isis csnp-interval** *seconds*
7. (Optional) **fabricpath isis hello-interval** *seconds*
8. (Optional) **fabricpath isis hello-multiplier** *multiplier*
9. (Optional) **fabricpath isis hello-padding**
10. (Optional) **fabricpath isis lsp-interval** *milliseconds*
11. (Optional) **fabricpath isis metric** *metric*
12. (Optional) **fabricpath isis retransmit-interval** *seconds*
13. (Optional) **fabricpath isis retransmit-throttle- interval** *milliseconds*
14. **exit**
15. (Optional) **show running-config**
16. (Optional) **copy running-config startup-config**

DETAILED STEPS

| | Command | Purpose |
|--------|---|--|
| Step 1 | <pre>configure terminal</pre> <p>Example: <pre>switch# configure terminal switch(config)#</pre></p> | Enters configuration mode. |
| Step 2 | <pre>interface {ethernet mod/slot port-channel channel-number}</pre> <p>Example: <pre>switch(config)# interface ethernet 5/2 switch(config-if)#</pre></p> | <p>Enters interface configuration mode and specifies the interfaces that you want to configure.</p> <p>The slot can be from 1 to 3. The following list defines the slots available:</p> <ul style="list-style-type: none"> • Slot 1 includes all the fixed ports. A Fabric Extender only has one slot. • Slot 2 includes the ports on the upper expansion module (if populated). • Slot 3 includes the ports on the lower expansion module (if populated). <p>The port number within a particular slot can be from 1 to 128.</p> <p>The port channel number assigned to the EtherChannel logical interface can be from 1 to 4096.</p> |
| Step 3 | <pre>fabricpath isis authentication-check</pre> <p>Example: <pre>switch(config-if)# fabricpath isis authentication-check switch(config-if)#</pre></p> | <p>(Optional) Enables authentication checking on incoming FabricPath Layer 2 IS-IS hello protocol data units (PDUs) for the interface. The default is ON. To disable authentication, enter the no form of the command.</p> <p>Note Level specification is not required.</p> |

| | Command | Purpose |
|--------|---|--|
| Step 4 | <p>fabricpath isis authentication key-chain <i>auth-key-chain-name</i></p> <p>Example: switch(config-if)# fabricpath isis authentication key-chain trees switch(config-if)#</p> | <p>(Optional) Assigns a password to authentication hello PDUs. To remove this password, enter the no form of the command.</p> <p>The maximum size for the name of an authentication keychain is 63 alphanumeric characters.</p> <p>Note Level specification is not required.</p> <p>An example of key chain creation is as follows:</p> <pre>key chain trees key 0 key-string cisco01 accept-lifetime 07:00:00 Sep 20 2011 infinite send-lifetime 07:00:00 Sep 20 2011 infinite</pre> <p>See the <i>Cisco Nexus 6000 Series NX-OS Security Configuration Guide, Release 6.0</i>, for information about key chains.</p> |
| Step 5 | <p>fabricpath isis authentication-type {cleartext md5}</p> <p>Example: switch(config-if)# fabricpath isis authentication-type md5 switch(config-if)#</p> | <p>(Optional) Specifies the authentication type for an interface for FabricPath Layer 2 IS-IS hello PDUs. To remove this type, enter the no form of the command.</p> <p>You can set one of the following authentication types:</p> <ul style="list-style-type: none"> • cleartext—Specifies the cleartext authentication method. • md5—Specifies the Message Digest (MD5) authentication. <p>Note Level specification is not required.</p> |
| Step 6 | <p>fabricpath isis csnp-interval <i>seconds</i></p> <p>Example: switch(config-if)# fabricpath isis csnp-interval 60 switch(config-if)#</p> | <p>(Optional) Specifies the interval in seconds between Complete Sequence Number (CSNP) PDUs sent on the interface. To return to the default value, enter the no form of this command.</p> <p>The range is from 1 to 65535, and the default value is 10.</p> |
| Step 7 | <p>fabricpath isis hello-interval <i>seconds</i></p> <p>Example: switch(config-if)# fabricpath isis hello-interval 20 switch(config-if)#</p> | <p>(Optional) Sets the hello interval between PDUs sent on the interface. To return to the default value, enter the no form of this command. The range is from 1 to 65535, and the default value is 10.</p> <p>Note Level specification is not required.</p> |

| | Command | Purpose |
|---------|---|---|
| Step 8 | <pre>fabricpath isis hello-multiplier multiplier</pre> <p>Example: <pre>switch(config-if)# fabricpath isis hello-multiplier 20 switch(config-if)#</pre></p> | <p>(Optional) Specifies the multiplier used to calculate the interval within which hello PDUs must be received or adjacency goes down. To return to the default value, enter the no form of this command. The range is from 3 to 1000. The default is 3.</p> <p>Note Level specification is not required.</p> |
| Step 9 | <pre>fabricpath isis hello-padding</pre> <p>Example: <pre>switch(config-if)# fabricpath hello-padding switch(config-if)#</pre></p> | <p>(Optional) Enables padding on the hello PDUs. The default is ON. To disable authentication, enter the no form of the command.</p> <p>Note If you enter the always keyword with the no form of this command, the padding is always on.</p> |
| Step 10 | <pre>fabricpath isis lsp-interval milliseconds</pre> <p>Example: <pre>switch(config-if)# fabricpath isis lsp-interval 100 switch(config-if)#</pre></p> | <p>(Optional) Sets the interval in milliseconds between link-state packets (LSPs) sent on this interface during flooding. To return to the default value, enter the no form of this command. The range is from 10 to 65535. The default is 33.</p> |
| Step 11 | <pre>fabricpath isis metric metric</pre> <p>Example: <pre>switch(config-if)# fabricpath isis metric 100 switch(config-if)#</pre></p> | <p>(Optional) Configures the FabricPath Layer 2 IS-IS metric for this interface. The range is from 0 to 16777215. To return to the default value, enter the no form of this command. The default values are as follows:</p> <ul style="list-style-type: none"> • 1 GB—400 • 10 GB—40 |
| Step 12 | <pre>fabricpath isis retransmit-interval seconds</pre> <p>Example: <pre>switch(config-if)# fabricpath isis retransmit-interval 100 switch(config-if)#</pre></p> | <p>(Optional) Sets the interval between initial LSP retransmissions. To return to the default value, enter the no form of this command. The range is from 1 to 65535. The default is 5.</p> |
| Step 13 | <pre>fabricpath isis retransmit-throttle- interval milliseconds</pre> <p>Example: <pre>switch(config-if)# fabricpath isis retransmit-throttle-interval 100 switch(config-if)#</pre></p> | <p>(Optional) Sets the interval between subsequent LSP retransmissions. To return to the default value, enter the no form of this command. The range is from 20 to 65535. The default is 66.</p> |
| Step 14 | <pre>exit</pre> <p>Example: <pre>switch(config-if)# exit switch(config)#</pre></p> | <p>Exits interface configuration mode.</p> |

| | Command | Purpose |
|---------|---|---|
| Step 15 | <pre>show running-config</pre> <p>Example: <pre>switch(config)# show running-config switch(config)#</pre></p> | (Optional) Displays the running configuration. |
| Step 16 | <pre>copy running-config startup-config</pre> <p>Example: <pre>switch(config)# copy running-config startup-config switch(config)#</pre></p> | (Optional) Copies the running configuration to the startup configuration. |

For more information about IS-IS commands, see the Unicast Routing configuration guide and command reference for your platform.

Clearing Advanced FabricPath Layer 2 IS-IS Counters

You can clear the FabricPath Layer 2 IS-IS counters.

SUMMARY STEPS

-
- Step 1 (Optional) **clear fabricpath isis adjacency** [* | **system-id** | **interface** {**ethernet** *mod/slot* | **port-channel** *channel-number*}]
 - Step 2 (Optional) **clear fabricpath isis statistics** *
 - Step 3 (Optional) **clear fabricpath isis traffic** [* | **interface** {**ethernet** *mod/slot* | **port-channel** *channel-number*}]

DETAILED STEPS

| | Command | Purpose |
|--------|--|---|
| Step 1 | <pre>clear fabricpath isis adjacency [* ethernet port-channel system-id {ethernet mod/slot port-channel channel-number system-id sid}]</pre> <p>Example: switch# clear fabricpath isis adjacency switch#</p> | <p>(Optional) Clears the FabricPath Layer 2 IS-IS adjacency state.</p> <ul style="list-style-type: none"> • *—Specifies the IS-IS adjacencies on all interfaces. • ethernet—Specifies the Ethernet interface. • mod/slot—The slot can be from 1 to 6. The port number within a particular slot can be from 1 to 96. • port-channel—Specifies the port channel interface. • channel-number—Port channel number. The range is from 1 to 4096. • system-id—Specifies the system ID. • sid—System ID in the form of XXXX.XXXX.XXXX. <p>Note If you enter the * variable, you affect forwarding which might interrupt traffic; this command tears down all adjacencies.</p> |
| Step 2 | <pre>clear fabricpath isis statistics *</pre> <p>Example: switch# clear fabricpath isis statistics * switch#</p> | <p>(Optional) Clears all FabricPath Layer 2 IS-IS protocol statistics.</p> |
| Step 3 | <pre>clear fabricpath isis traffic { * ethernet mod/slot [. sub-int] port-channel channel-number }</pre> <p>Example: switch# clear fabricpath traffic switch#</p> | <p>(Optional) Clears FabricPath Layer 2 IS-IS traffic information.</p> <ul style="list-style-type: none"> • *—Specifies the IS-IS adjacencies on all interfaces. • ethernet—Specifies the Ethernet interface. • mod/slot—The slot can be from 1 to 3. The port number within a particular slot can be from 1 to 128. • sub-int—Specifies that sub-interface number. • port-channel—Specifies the port channel interface. • channel-number—Port channel number. The range is from 1 to 4096. |

Multi-Destination Tree

Information About MDT

The Multi-Destination Tree (MDT) feature removes the restriction of two MDTs per topology and permits a user to configure up to 16 MDTs per topology. This feature allows load-sharing of multicast traffic across different multi-destination trees within a topology, thus improving the network performance.

Guidelines and Limitations for MDT

An MDT, also referred to as a Forwarding Tag or ftag, is a spanning-tree used for forwarding packets within a topology. By default, a topology has two MDTs ftags: topology 0 has ftag 1 and 2, topology 1 has ftag 3 and 4, topology 2 has ftag 5 and 6. In a default configuration, the first ftag of a topology is used for broadcast and unknown unicast. The second ftag is used for multicast only. In addition, multicast traffic might use the first ftag for load-sharing.



Note

The 256 and 257 ftags are reserved ftags used for a CE VPC topology.

Configuration and usage of MDT on the Cisco Nexus device needs to follow the following rules:

- **Minimum of Maximum MDTs:** A contiguous fabricpath network has the maximum number of multi-destination trees (MDT) configured on every node in the network. Each node advertises the maximum number of trees it can support. The minimum value of the max trees across the nodes is used as the final operational count for MDT.

For example, if there are ten switches in a fabricpath network, and if nine switches are configured with 8 MDTs but the tenth switch is not configured with MDT, then by this rule, the tenth switch advertises its maximum supported trees as two. Each switch has two MDTs which is the minimum number of MDTs in the fabricpath network.

- **Number of zero rooted trees with regards to the number of nodes:** The number of multi-destination trees with zero root is equal to the number of MDTs configured minus the number of nodes in the fabricpath network.

For example, if there are three switches in the fabricpath network and each of these switches is configured with 8 MDTs, each switch would have $(8 - 3 = 5)$ zero rooted trees. Only three trees are used with each node acting as the root for each tree.

- **Maximum number of MDTs per topology:** The maximum number of MDTs supported per topology is 16. If no MDT is configured on a topology, the default of two MDTs takes effect.
- **Total number of MDTs across all topologies:** The maximum number of MDTs supported functionally is 16 across all topologies. For example, if topology-0 has 8 trees, topology-1 is allowed to have a maximum of eight trees functionally.



Note

The configuration of 16 MDTs for 16 topologies passes through, However, functionally we plan to support maximum 16 MDTs across all topologies

- You can configure more than two MDTs for topologies 0 through 15.

- **MDT Root Unreachable:** Root re-calculation for the trees is triggered whenever any node in the setup becomes unreachable. This can also affect traffic in the network. For example, if there are four nodes with four MDTs configured with root priority as Node1>Node2>Node3>Node4 and Node 2 (Root for tree 2) reboots, root re-calculation is triggered as per the node priority. In this case, Node 3 becomes root for Tree2, Node 4 becomes root for Tree 3 and Tree 4 will be zero-rooted. Traffic for all the trees whose path went through Node 2 is affected temporarily in this case. The higher the priority of the disabled root, the more traffic is disrupted.
- **Number of nodes is less than the number of MDTs:** If the MDT configuration changes during run-time such that the number of MDTs configured is greater than the number of nodes in the network, the root-calculation of traffic is not affected. This adheres to the minimum of maximum rule.

For example, if there are four nodes and four MDT configured and later the admin configures eight MDTs on all four nodes, we see four MDTs and the root calculation remains the same.

- **Mismatch in MDT Configuration in the beginning:** If each switch in the fabricpath network has a different number configured as the number of MDTs, again the minimum of maximum rule applies and we see the least number of trees on all nodes.

For example, if there are three switches in the fabricpath network and switch 1 = 3 MDTs, switch 2 = 5 MDTs, switch 3 = 8 MDTs, the total number of MDTs seen on the setup is three.

Configuring MDT for Default Topology

You can configure MDT for a default topology.

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# configure terminal | Enters global configuration mode. |
| Step 2 | switch(config)# fabricpath domain default | Enters fabricpath IS-IS domain configuration mode. |
| Step 3 | switch(config-fabricpath-isis)# multi-destination trees tree-num | Specifies the number of multi-destination trees for a topology. |

Configuring Per-Topology MDT

You can configure MDT per topology.

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# configure terminal | Enters global configuration mode. |
| Step 2 | switch(config)# fabricpath domain default | Enters fabricpath IS-IS domain configuration mode. |
| Step 3 | switch(config-fabricpath-isis)# topology topology number | Specifies the Layer 2 fabricpath topology ID. The range of <i>topology number</i> is from 1 to 63. |
| Step 4 | switch(config-fabricpath-isis-topo)# multi-destination trees tree-num | Specifies the number of multi-destination trees for a topology. The default value is 2. |

Verifying the FabricPath Advanced Configurations

To display FabricPath information for advanced configurations perform one of the following tasks:

| Command | Purpose |
|--|---|
| show fabricpath isis adjacency [{ ethernet <i>mod/slot</i> port-channel <i>channel-number</i> } system-id detail summary] | Displays the FabricPath Layer 2 IS-IS adjacency database. |
| show fabricpath isis database [level-1 [mgroup] [detail summary] { zero-seq router-id adjacency }[<i>SID.XX-XX</i>] | Displays the FabricPath Layer 2 IS-IS database. |
| show fabricpath isis hostname [detail] | Displays the FabricPath Layer 2 IS-IS dynamic hostname exchange information. |
| show fabricpath isis interface [ethernet <i>mod/slot</i> port-channel <i>channel-number</i>] [brief] | Displays the FabricPath Layer 2 IS-IS related interface information. |
| show fabricpath isis route [summary detail] | Displays the FabricPath Layer 2 IS-IS routing table for unicast routes. |
| show fabricpath isis spf-log [detail] | Displays the FabricPath Layer 2 IS-IS SPF calculation statistics. |
| show fabricpath isis statistics | Displays the FabricPath Layer 2 IS-IS event counters. |
| show fabricpath isis ftag [multidestination <i>tree_id</i>] | Displays the FTag values associated with the trees in the topology. |
| show fabricpath isis vlan-range | Displays the congruent VLAN-set to topology mapping. |
| show fabricpath isis trees [multidestination <i>tree_id</i>] | Displays the nodes in the trees, reachability to these nodes through a specific interface, and the configured metric. |
| show fabricpath isis switch-id | Displays the switch IDs and reachability information for the topology. |
| show fabricpath isis ip redistribute mroute [vlan <i>vlan-id</i> [group <i>group-id</i> [source <i>source-id</i>]]] | Displays the locally learned multicast routes. |
| show fabricpath isis ip mroute [vlan <i>vlan-id</i> [group <i>group-id</i> [source <i>source-id</i>]]] | Displays the multicast routes learned from neighbors. |
| show fabricpath isis [protocol] | Displays the FabricPath Layer 2 IS-IS process level information. |
| show fabricpath isis rrm [gm] { ethernet <i>mod/slot</i> port-channel <i>channel-number</i> } | Displays the FabricPath Layer 2 IS-IS retransmit-routing-message information. |
| show fabricpath isis srm [gm] { ethernet <i>mod/slot</i> port-channel <i>channel-number</i> } | Displays the FabricPath Layer 2 IS-IS send-routing-message information. |
| show fabricpath isis topology summary | Displays the FabricPath Layer 2 IS-IS topology database. |

| Command | Purpose |
|---|---|
| <code>show fabricpath isis traffic [{ethernet <i>mod/slot</i> port-channel <i>channel-number</i>}]</code> | Displays the FabricPath Layer 2 IS-IS traffic information. |
| <code>show fabricpath isis ssn [gm] {ethernet <i>mod/slot</i> port-channel <i>channel-number</i>}</code> | Displays the FabricPath Layer 2 IS-IS send-sequence-number information. |

Feature History for Configuring FabricPath Advanced Features

Table 5-1 lists the release history for these features.

Table 5-1 Feature History for FabricPath

| Feature Name | Releases | Feature Information |
|------------------------------|-------------|--|
| Advanced FabricPath features | 6.0(2)N1(1) | These features were introduced with support for IS-IS. |



Configuring FabricPath in Transit Mode

This chapter contains the following sections:

- [“Information about FabricPath in Transit Mode” section on page 6-1](#)
- [“Guidelines and Limitations” section on page 6-1](#)
- [“Configuring FabricPath in Transit Mode” section on page 6-2](#)
- [“Feature History for Configuring FabricPath in Transit Mode” section on page 6-3](#)

Information about FabricPath in Transit Mode

In transit mode, all the incoming traffic is mapped to one internal bridge-domain (VLAN). A pure Layer 2 spine node does not learn the MAC addresses; however, the packet is switched based on the switch ID lookup table in the FabricPath header.

In a FabricPath network, the Layer 2 spine node can be configured in Transit mode. After the spine is configured in the Transit mode, uni-destination traffic received on the spine is forwarded based on the FTAG (forwarding tag) and destination switch ID. For multi-destination traffic, the traffic on the spine is forwarded based on FTAG, which is the identifier of the multi-destination traffic that is selected by the ingress switch.

The spine will forward all the fabric-path data for all VLANs. For any other traffic, for example, carrier ethernet ports or SVI or Layer 3-specific data, processing is dependent on the SVI or Layer 3 configuration, based on SVI enabled on the spine switch.

Guidelines and Limitations

The following are the guidelines and limitations for configuring FabricPath in Transit mode:

- Ensure that a minimum of one VLAN is enabled in FabricPath mode. The default VLAN “VLAN1” can be enabled as FabricPath mode. However, we recommend you to set another VLAN as FabricPath mode.
- In case a SVI is enabled on spine node, the SVI is only intended to be used as management purpose. It should not carry data traffic because data traffic in that SVI VLAN may not be forwarded correctly on transit mode spine.

Configuring FabricPath in Transit Mode

To configure the FabricPath in transit mode, perform the following.

Procedure

-
- Step 1** Enter the global configuration mode:
- ```
switch# configure terminal
switch(config)#
```
- Step 2** Install the FabricPath feature set on the switch.
- ```
switch(config)# install feature-set fabricpath
```
- Step 3** Enables the FabricPath feature set on the switch.
- ```
switch(config)# feature-set fabricpath
```
- Step 4** Configure the FabricPath in transit mode:
- ```
switch(config)# fabricpath mode transit
```
- Step 5** Copy and save the running configuration to the startup configuration.
- ```
switch(config)# copy running-config startup-config
```
- Step 6** Reload the spine for transit mode to take effect.
- ```
switch(config)# reload
```
-

Example: Configuring FabricPath in Transit Mode

This example shows a running configuration, followed by a verification command that displays the configuration details for Cisco FabricPath in Transit Mode.

```
switch(config)# fabricpath mode transit
Enabling transit mode. Please save configuration and reload

switch(config)# show fabricpath mode
Switch is configured in transit mode. Please save configuration and reload to run in
transit mode

switch (config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...

switch (config)# reload
```

The following is the output for the **show fabricpath mode** command.

```
switch# show fabric path mode
Switch is configured and running in transit mode.
```


Feature History for Configuring FabricPath in Transit Mode

The following table lists the release history for these features.

Table 6-1 Feature History for Configuring FabricPath in Transit Mode for Nexus 6000 switches

| Feature Name | Releases | Feature Information |
|----------------------------|-------------|------------------------------|
| FabricPath in Transit Mode | 7.0(0)N1(1) | This feature was introduced. |



Configuration Limits for Cisco NX-OS FabricPath

For the maximum configuration limits for the Cisco NX-OS FabricPath features supported by your platform, see the Configuration Limits document document at this URL:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html



INDEX

A

anycast HSRP [127](#)
 configuring anycast bundle [135](#)

B

broadcast traffic [144](#)

C

CE VLANs [18](#), [123](#), [142](#), [145](#), [146](#)
configuring [145](#), [146](#)
conflict resolution [110](#)
 displaying [120](#)
conversational learning [18](#), [19](#)

D

DHCP snooping [14](#)
disabling [114](#)
documentation
 related documents [16](#)

E

ECMP [141](#), [143](#)
enabling [14](#), [111](#)
 configuring [113](#)
 example [138](#)
encapsulation
 EtherType [17](#)
 FTag [17](#)

F

FabricPath
 configuring vPC+ [134](#)
 switch ID [110](#)
 timers [110](#)
FabricPath forwarding [141](#)
 unicast traffic [143](#)
FabricPath VLANs [18](#), [123](#), [142](#)
 configuring [145](#), [146](#)
 example configuration [150](#)
feature set
 displaying [137](#)
forwarding
 configuring [145](#)
 displaying [149](#)
 example configuration [150](#)
 multidestination traffic [144](#)
 trees [141](#)
FTag [17](#)
FTags [17](#), [110](#), [141](#), [144](#)

G

GM-LSP [144](#)
graceful merge
 configuring [118](#)
graceful migration [110](#)
guidelines [14](#)

I

IGMP Snooping [145](#)

PIM 145

interfaces

clearing counters 138

displaying 137

examples 138

network 124

statistics 138

STP 124

tunnel 134, 136

IS-IS overload bit

configuring 154

description 154

K

known unicast traffic 143

forwarding 18

L

Layer 2 gateway switch 124

STP priority 130

layer2 gateway switches 133

Layer 2 IS-IS 17

adjacencies 17, 143

advanced settings 153

clearing 163

configuring 154

displaying 167

ECMP 17

global parameters 154

interface parameters 159

trees 142

limitations 14

load balancing

displaying 149

unicast traffic 148

M

MAC address

hierarchical 17

MAC addresses 19

conversational learning 18

mobility 19

multicast traffic 144

groups 144

IGMP Snooping 145

multidestination packets 142

multidestination traffic 144

forwarding 18

multipathing at Layer 2 141

multiple paths

FTag 17

O

ODA 16

OSA 16

P

ports 123

R

RPF 142

S

Spanning Tree Protocol 17

STP 124, 125

displaying 137

root 124

separate connected domains 133

TCNs 125

STP domain [138](#)
 STP domains [133](#)
 STP priority [130](#)
 example [138](#)
 switch ID [16, 110](#)
 conflict [110](#)
 displaying [120](#)
 switching
 clearing counters [120](#)
 default settings [121](#)
 example [19](#)
 prerequisites [13](#)
 statistics [120](#)
 verifying [120](#)

T

TCNs [125, 133](#)
 timers [110](#)
 allocate-delay [117](#)
 configuring [116, 117](#)
 displaying [120](#)
 linkup-delay [117](#)
 transition-delay [117](#)
 TLV [110](#)
 TLV information [18](#)
 topology
 FabricPath [143](#)
 traditional learning [18](#)
 traffic
 CE [143](#)
 FabricPath [143](#)
 trees [141, 144](#)
 displaying [149](#)
 troubleshooting [18, 110, 119, 125, 126, 143, 145](#)
 enabling [14, 111](#)
 traffic drops [118](#)
 trunk ports [123](#)
 TTL [17](#)

tunnels
 creating [134, 136](#)

U

unicast traffic [143](#)
 unknown unicast traffic [144](#)

V

VLAN modes [18](#)
 VLANs [18, 123, 145, 146](#)
 default [151](#)
 default mode [143](#)
 displaying [149](#)
 example configuration [150](#)
 FabricPath [142](#)
 troubleshooting [143](#)
 [143](#)
 vPC+ [125, 126](#)
 configuring [139](#)
 displaying [137](#)
 switch ID [134](#)

